

# Contents

<b>Abstract</b> .....	<b>iii</b>
<b>Preface</b> .....	<b>vii</b>
Acknowledgements .....	vii
How to read this thesis .....	viii
<b>Contents</b> .....	<b>xi</b>
<b>List of figures and tables</b> .....	<b>xvii</b>
List of figures .....	xvii
List of tables and boxes .....	xix
<b>Glossary of terms</b> .....	<b>xxiii</b>
<b>Chapter 1 Introduction</b> .....	<b>1</b>
Origins and drivers of ERM and the CRO role .....	4
Drivers of ERM uptake .....	5
Continuing interest in ERM and the CRO role .....	6
The state of the CRO profession .....	7
The state of formal representation of CROs .....	8
The state of knowledge base support for the CRO role .....	9
Where and how this thesis contributes .....	16
A Theory of ERM: explaining the effects of ERM implementation .....	17
A Theory of Action for the CRO: explaining how to perform the role.....	19
The relationship between the two theories .....	20
Locating the contributions of this thesis .....	22
Contributions in engineering.....	25
Organisation of the thesis.....	28
<b>Chapter 2 Methodology</b> .....	<b>31</b>
On performing reliable knowledge .....	32
There is nothing so practical as a good theory .....	34
From correspondence to resistance: the objectivity of strong networks .....	36
The “pragmatic turn” in the social sciences .....	40
So we should all be pragmatists?.....	42
Performing the role of Chief Risk Officer: a pragmatic inquiry .....	44
Shadowing as a way to fulfil a practical interest (and mitigate risk) .....	46
Two projects: collaboratively seeking understanding of practice .....	49
Pragmatic hermeneutics in response to strategic uncertainty .....	53
Judging the account: workability .....	61

---

Judging the account: warrantability .....	62
On how I became a transdisciplinary .....	69
<b>Chapter 3 The CRM's vision for what ERM should look like .....</b>	<b>73</b>
The CRM's proposals for developing ERM at Watercare .....	75
Concerns about Watercare's "risk data" .....	75
Concerns about the existing risk framework (2003) .....	76
Concerns about risk communication .....	77
The evidence from the CRM's survey of staff perceptions .....	80
A vision for better "risk data" .....	81
Comparing the CRM's proposals against normative models .....	85
The CRM's rationale: supporting and defending decisions with "good risk data" .....	89
Discussion: ERM and the construction of rational management in organisations .....	92
<b>Chapter 4 Experiences with designing ERM infrastructure .....</b>	<b>97</b>
Redefining Watercare's risk assessment framework .....	98
The design rationale .....	98
The art of marrying abstract objectives to concrete practices .....	100
Dilemmas .....	106
Difficulties using the new Project risk register .....	108
Specifying risk requires a certain degree of knowledge .....	108
Knowledge of risk has to be produced .....	110
The risk register as a technology of bureaucracy .....	111
Limits to modelling the Enterprise Risk Profile .....	114
Facilitating the identification of strategic risks .....	115
Sense-making in a context of uncertainty, ambiguity, and vagueness .....	116
Performing the role of knowledge facilitator .....	119
Discussion .....	123
Ontological lessons .....	125
Epistemological lessons .....	130
The value problem: how much detail? .....	137
<b>Chapter 5 Rethinking the CRM's approach to decision support .....</b>	<b>141</b>
Rethinking the role of "risk data" (Dialogues 11 & 12) .....	142
Rethinking Risk Management (Dialogue 13) .....	148
The CRM rethinks his approach (Dialogue 14) .....	155
Discussion: "strategic controller" in question .....	158
Reinterpreting Risk Management .....	161
The CRO's dilemma: how to convince actors to adopt more advanced capabilities ..	168
The potential to add value as a "strategic advisor" .....	169
How should CROs support organisational decision making? .....	171
<b>Chapter 6 Framing up decision support strategies for Chief Risk Officers..</b>	<b>173</b>

---

The status of rational choice theory.....	176
Understanding real decision making behaviour.....	180
Case 1: the formalised construction of a rational choice .....	180
Case 2: putting on a skillful performance .....	183
Factors affecting framing and calculative processes .....	186
The relationship between agent and context.....	190
A typology of problem/decision situations.....	193
Extending the typology .....	197
Reconceptualising “decision” in performative terms .....	205
Framing as constitutive of decision .....	205
Sources of uncertainty in decision frames.....	207
Conditions for cutting off framing reasonably.....	209
Decision support strategies for Chief Risk Officers.....	210
Strategy 1: Formatting the rational decision maker.....	213
Strategy 2: Providing resources to the decision maker .....	215
Strategy 3: Helping decision makers to settle their frames .....	215
<b>Chapter 7 Conclusions .....</b>	<b>227</b>
Conclusions .....	228
Contributions.....	234
Future Research.....	235
<b>Dialogues .....</b>	<b>237</b>
Dialogue No. 1 – The Vision .....	239
Dialogue No. 2 – Quantifying risk profiles .....	247
Dialogue No. 3 – Proposed developments .....	255
Dialogue No. 4 – Engaging with stakeholders: Operations Managers.....	261
Dialogue No. 5 – Engaging with stakeholders: Project Managers.....	273
Dialogue No. 6 – Engaging with stakeholders: Asset Managers.....	283
Dialogue No. 7 – Engaging with stakeholders: General Managers .....	291
Dialogue No. 8 – Reflecting on the dialogue with the General Managers .....	297
Dialogue No. 9 – Justifying the risk framework .....	303
Dialogue No. 10 – Feedback on strategic risks .....	307
Dialogue No. 11 – The purpose of “risk data” .....	313
Dialogue No. 12 – On the value of the risk framework in a mature organisation.....	319
Dialogue No. 13 – Trying to explain the role.....	325
Dialogue No. 14 – Reflections on a year in the role .....	331
<b>Appendices .....</b>	<b>337</b>
<b>Appendix I The research and practice context .....</b>	<b>339</b>

---

The corporate context .....	339
Infrastructure, services and performance standards .....	340
Corporate status and governance .....	343
Mission and objectives.....	345
Company structure.....	348
Engineering asset planning and management.....	349
Collaborative regional planning.....	353
Capital decisions: the Service-Cost-Risk trade-off .....	354
The risk assessment framework .....	355
Scoring and classification .....	356
Use of the risk framework.....	357
Significance and quality of the Risk Management function.....	359
Key performance statistics for the period 2000 - 2008.....	361
Value-For-Money .....	362
Non-compliance with regulatory standards for service delivery .....	362
Asset management processes .....	364
Other.....	365
Summary of Watercare's performance 2000-2008 .....	366
<b>Appendix II Early experiences crossing disciplinary boundaries.....</b>	<b>369</b>
Early work: attempting to engineer sustainability .....	369
Redesigning decision making infrastructure .....	374
Redesigning Watercare's risk assessment framework .....	375
Redesigning Watercare's capital planning framework .....	382
Extending the Project Delivery Manual.....	387
More transdisciplinary work: reframing the inquiry .....	388
The evolution of Large Technical Systems .....	389
Characteristics of change in Auckland's water infrastructure.....	396
General systems management model of the Watercare organisation .....	399
The nature of engineering planning and design .....	401
Identifying potential contexts for inquiry.....	405
Recognising the opportunity presented by the CRM's project.....	409
<b>Appendix III Review of literature around ERM and the CRO role.....</b>	<b>411</b>
Drivers of ERM uptake .....	411
Corporate interest in ERM and the CRO role.....	417
Professional representation of Chief Risk Officers .....	422
Review of practical literature on the implementation of ERM and the CRO role.....	423
Review of academic literature on the implementation of ERM and the CRO role.....	428
On Enterprise Risk Management .....	428
On the role of Chief Risk Officer.....	431
<b>Appendix IV Proposals to improve the quality of Watercare's "risk data" ...</b>	<b>437</b>

Concerns about Watercare's "risk data" .....	437
Proposals to improve data quality .....	440
<b>Appendix V Normative models of Risk Management .....</b>	<b>445</b>
The principles of ERM.....	445
The distinguishing features of ERM.....	446
Cross-functional integration – the portfolio view.....	447
Embedding risk management – creating the risk culture .....	448
Centering the corporate objectives.....	449
Normative models for benchmarking ERM .....	449
Risk assessment in engineering contexts .....	452
Why a functional systems perspective?.....	453
Engineering ideals for risk assessment .....	454
Defining context .....	456
Identifying risk .....	459
Analysing risk .....	460
<b>Appendix VI Experiences with implementing ERM.....</b>	<b>465</b>
(Re)defining Watercare's risk assessment framework .....	465
The design task.....	466
Engaging with stakeholders.....	468
The revised (2008) Operations Risk Framework.....	480
Redesigning the Project Risk Register .....	483
The new Project Risk Register .....	483
Surprising difficulties using the new register .....	486
The availability and quality of data for risk modelling.....	488
Facilitating the identification of strategic risks.....	490
Identifying strategic risks.....	491
Defining the term "strategic risk" in Watercare's context .....	493
Rationalising and categorising strategic uncertainties .....	496
Fleshing out the details of Watercare's strategic risks .....	498
<b>Appendix VII Examining the CRM's contradiction .....</b>	<b>501</b>
The implications of View 1.....	503
The implications of View 2.....	503
Resolving the contradiction.....	504
<b>Appendix VIII Decision: cutting off the framing process.....</b>	<b>511</b>
On moving from knowing to acting .....	511
How does belief operate? .....	512
Is truth necessary for rational action? .....	515
What constitutes justified belief?.....	515
What constitutes reliable evidence? .....	516
What constitutes sufficient evidence for rational belief?.....	516

What are the normative conditions for rational action in the face of incomplete knowledge? .....	518
Conditions for settling decision frames reasonably .....	520
Resolving the satisficing conundrum .....	521
Making the satisficing cut-off (specific case) .....	522
General formulation of the satisficing decision .....	530
<b>References .....</b>	<b>535</b>

# List of figures and tables

Figures and tables are numbered consecutively according to the Chapter or Appendix in which they are located. Thus:

Figure 1.1 = Figure 1 in Chapter 1

Table 6.5 = Table 5 in Chapter 6

Figure II.2 = Figure 2 in Appendix II

Table III.3 = Table 3 in Appendix III

## List of figures

<b>Figure 1.1.</b> Current state of the risk management knowledge base .....	11
<b>Figure 1.2.</b> Conceptual model relating a Theory of Enterprise Risk Management and a Theory of Action for the role of Chief Risk Officer. ....	21
<b>Figure 2.1.</b> The Hermeneutic Circle.....	59
<b>Figure 2.2.</b> Key work streams and questions constituting the inquiry .....	60
<b>Figure 2.3.</b> Presentation of the thesis.....	63
<b>Figure 4.1.</b> Illustration of the CRM's rationale for how the structure and content of risk frameworks can affect the representation and communication of risk between communities of practice in the organisation: the proper alignment of risk framework categories to corporate objectives and performance standards, and the proper specification of risk evaluation criteria in common operational terms, are assumed to improve the objectivity of risk communication. ....	101
<b>Figure 4.2.</b> The distinction between Risk and Uncertainty (modified from Rosa 1998, Figure 3, p 35).....	134
<b>Figure 4.3.</b> The assumed knowledge production contexts of the CRM's interventions discussed in this chapter. The dark shaded arrow represents states of knowledge about the world as a progression from uncertainty to risk (cf. Figure 4.2) .....	134
<b>Figure 6.1.</b> The relationship between agent and context with respect to framing. The x-axis represents the problem situation, as it is perceived by the agent, and the y-axis the nature of the problem solving process, as it is experienced by the agent. Both axes are therefore subjective to the agent. ....	191
<b>Figure 6.2.</b> Different uncertainties experienced by Watercare's planners and the race car driver .....	193

<b>Figure 6.3.</b> Problem situations defined in terms of the agent-context relationship. The x-axis represents the problem situation, as it is perceived by the agent, and the y-axis the nature of the problem framing / solving process, as it is experienced by the agent. The axes, and the problem definitions, are relative (i.e. subjective) to the agent. ....	196
<b>Figure I.1.</b> Former structure of the Auckland water industry (compiled from Watercare Services Ltd Annual Reports 2002, 2004, & 2007) .....	344
<b>Figure I.2.</b> Watercare’s general infrastructure planning and development process.....	351
<b>Figure I.3.</b> The general development and operational life cycle of infrastructure assets. Numbers indicate key locations for use of Watercare’s risk assessment framework. Figure originally developed for draft of Watercare’s Project Development & Delivery Manual 2007 .....	352
<b>Figure I.4.</b> The formal policy, scope, and objectives for Risk Management at Watercare (source: WSL Risk Management Framework manual 2003, p 3).....	356
<b>Figure I.5.</b> Public commentary on risk management in WSL Annual Reports 2003 - 2005 (sources: Watercare Services Ltd. 2003a, 2004a; 2005a, page numbers as indicated).....	360
<b>Figure I.6.</b> Watercare’s wholesale prices for water and wastewater services 2000 – 2008. Source: calculated from annual turnover versus volumes supplied/treated quoted in Watercare Annual Reports 2000 – 2008; cf. also Figure 63, p 144 of the 2008 AR. CPI adjustments calculated from CPI figures obtained from Statistics NZ ( <a href="http://www.stats.govt.nz">www.stats.govt.nz</a> ). .....	362
<b>Figure II.1.</b> Indicative conceptual map of my exploration of “sustainability” and related concepts. Arrows indicate linkages between concepts, numbers indicate density of connections to key subject areas in the literature. ....	373
<b>Figure II.2.</b> Generic concept for translating Watercare’s corporate objectives into a comprehensive performance evaluation framework (adapted from presentation to managers, 26 May 2005). ....	384
<b>Figure II.3.</b> Conceptual scope of integrated performance evaluation framework and relationships to key organisational processes (source: PEF, Proposed Solution, 29 November 2006) .....	385
<b>Figure II.4.</b> Indicative mapping of actors and forces actually or potentially influencing the development and delivery of water and wastewater services in the Auckland region. ....	397
<b>Figure II.5.</b> Generic management systems model of organisations, after Beer’s Viable Systems Model (Beer 1972, 1985; Espejo and Harnden 1989) .....	401
<b>Figure II.6.</b> Possibilities and limitations on the state of technology .....	409
<b>Figure V.1.</b> Generic relational structure of statements of risk .....	460
<b>Figure VI.1.</b> Different objectives for different contexts (source: Clement 2007d) .....	467
<b>Figure VI.2.</b> General systems framework for categorising strategic sources of uncertainty .....	497



## List of tables and boxes

Table 1.2. General assumptions about the effects of ERM implementation (in the world-level literature).....	18
Table 2.1. Parallel streams of inquiry .....	50
Box 2.1. Relevance of different literatures to understanding Risk and Risk Management .....	51
Table 2.2. Evolving assumptions about key objects of attention.....	68
Table 2.3. Prior experiences prepared me to follow Watercare's Corporate Risk Manager .....	71
Table 3.1. CRM's concerns about the quality of "risk data" in Watercare's corporate risk register .....	75
Table 3.2. Comparing the CRM's vision against normative ERM principles.....	86
Table 3.3. Maturing Watercare's Risk Management (RM) capabilities .....	87
Table 3.4. The CRM's desired improvements for "risk data" quality .....	89
Table 4.1. The CRM's definitions of key risk categories (source: Clement 2008e, Fig. 2, p 7) .	118
Table 4.2. Evolution of the strategic risk framework and template .....	123
Table 4.3. The CRM's score card.....	124
Table 6.1. Comparing two decision contexts and performances .....	187
Box 6.1. Factors influencing the ability of actors to frame and calculate in context .....	190
Table 6.2. A basic typology of problem/decision situations (see also Figure 6.5) .....	195
Table 6.3. Nature of epistemological and axiological uncertainties defining problem domains (Extends Table 6.2 by reference to Funtowicz and Ravetz 1993) .....	198
Table 6.4. System roles in Funtowicz and Ravetz (1993) .....	199
Table 6.5. Directive and Evaluative functions of problem-solver and stakeholder roles .....	201
Table 6.6. Framework characterising the general features of problem domains, the problem solver roles associated within those domains, and the functions of those roles .....	203
Table 6.7. Types and sources of epistemic uncertainty in decision frames .....	209
Table 6.8. A summary of decision making .....	212
Table 6.9. Two paradigmatic views of the Chief Risk Officer's role in decision support .....	219
Table 6.10. Business processes address different types of problems .....	221
Table 6.11. Conceptualising Strategy 3 as planning support in different contexts within the firm .....	222

<b>Table I.1.</b> Summary of Watercare’s water and wastewater infrastructure, and budgets .....	342
<b>Table I.2.</b> Watercare’s corporate objectives as specified in the SCI and Annual Report .....	347
<b>Table I.3.</b> Dimensions of Watercare’s performance and corresponding key performance indicators.....	361
<b>Table I.4.</b> Watercare performance: Non-compliance with regulatory standards 2000-2008.....	363
<b>Table I.5.</b> Watercare performance: Variance against Operational and Capital Expenditure Budgets.....	364
<b>Table I.6.</b> Watercare performance: Miscellaneous 2000-2008 .....	365
<b>Table I.7.</b> Summary Watercare’s performance from 2000-2008.....	367
<b>Table II.1.</b> Different usages of the “sustainability” label.....	374
<b>Table II.2.</b> Revised Risk Assessment Framework that I developed with Watercare’s Risk Specialist in 2005 .....	379
<b>Table III.1.</b> Selected corporate governance, control, and risk management initiatives .....	413
<b>Table III.2.</b> Comparative summary of surveys on ERM uptake.....	420
<b>Table III.3.</b> Comparative summary of industry surveys showing penetration of the role of CRO amongst global corporates.....	421
<b>Table III.4.</b> Sources for Tables III.5 and III.6.....	425
<b>Table III.5.</b> Common functional aspects of the role of CRO (sources for this table are listed in Table III.4, above).....	426
<b>Table III.6.</b> Attributes and skill sets of effective CROs (sources for this table are listed in Table III.4, above).....	427
<b>Table III.7.</b> Functions and objectives of the CRO role, and factors influencing the direction and scope of the role (from Ward 2001) .....	431
<b>Table IV.1.</b> The CRM’s proposals to improve data quality (collated by function).....	441
<b>Table V.1.</b> Characteristics of ERM, as defined in world-level frameworks.....	446
<b>Table V.2.</b> Comparison of selected Capability Maturity Models for Enterprise Risk Management (references in text on previous page) .....	451
<b>Table V.3.</b> Typical adjectives used to describe risk management process maturity.....	452
<b>Table V.4.</b> Contextual elements of risk.....	458
<b>Table VI.1.</b> Measuring performance against corporate objectives .....	481
<b>Table VI.2.</b> Comparison of the evaluation categories in the existing 2003 and revised 2007 project risk registers.....	485

**Table VII.1.** Contradiction between the CRM's expressed views of the Watercare organisation .....502

**Table VIII.1.** Basis for belief in P1 and P2.....526

**Table VIII.2.** Basis for belief in P3 .....529

**Table VIII.3.** Basis for belief in P1, P2, and P3 .....534



# Glossary of terms

*Common abbreviations used in this thesis are as follows:*

AMP	Asset Management Plan
Capex	Capital Expenditure Request
CMM	Capability Maturity Model
CRM	Corporate Risk Manager (unless otherwise indicated refers to the specific role at Watercare Services Ltd)
CRO	Chief Risk Officer (refers to the role of CRO in general)
ERM	Enterprise Risk Management
RCM	Reliability Centred Maintenance
RM	Risk Management
SCI	Statement of Corporate Intent (unless otherwise indicated refers specifically to Watercare's SCI)
WSL	Watercare Services Ltd

*Quotes and Quotations:*

- `Quotes from the Dialogues are identified by this font`
- `Quotations from documentary sources are identified by this font`



# Chapter 1

## Introduction

This thesis is about the role of the Chief Risk Officer (CRO) and the implementation of Enterprise Risk Management (ERM) in organisations. The CRO is that corporate role responsible for the leadership, co-ordination, and integration of risk management across the enterprise (Conference Board of Canada 2001; Lam and Kawamoto 1997; Power 2005b; Ward 2001). The role is referred to by a variety of other titles, including Corporate Risk Manager (Ward 2001) and is usually taken as a sign that an organisation has adopted an ERM programme (Beasley, Clune, and Hermanson 2005a; Liebenberg and Hoyt 2003; Power 2005b).

Enterprise Risk Management does not refer to a unified set of practices or a single standard, but rather to a category of standardised and widely adopted designs and process frameworks for the institutionalisation of rational comprehensive risk management in organisations (Power 2005a, 2007). The concept, processes, and components of ERM are defined in a number of international standards, frameworks, and guideline documents, the most prominent of which are the COSO *Enterprise Risk Management – Integrated Framework* (2004), the Australia/New Zealand Standard *4360 Risk Management* (2004a), and the more recent *ISO 31000 Risk Management – Principles and Guidelines* (2009). Across these documents, the common elements of a definition of ERM are: an approach to the identification, analysis, evaluation, and management of risk which is structured, rigorous, and systematic, with an holistic, enterprise-wide focus, and seamlessly integrated into all aspects of organisational decision making such that risk management becomes part of the organisation culture (see Appendix V).

The thesis is grounded by a close following of the Corporate Risk Manager (CRM) at Watercare Services Ltd (Watercare or WSL) between February 2007 and May 2008 as he

sought to implement an ERM programme within the company<sup>1</sup>. From here on, the term “Chief Risk Officer” or “CRO” will be used to refer to the role in general, while the term “Corporate Risk Manager” or “CRM” will be used to refer specifically to the CRM’s role at Watercare. The thesis focuses on the significant shift in approach which characterised the CRM’s performance:

- From an initial perception of the role framed by a normative focus on the objectivity and precision of “risk data”, grounded in a conceptual understanding of risk as something objectively calculable and specifiable, and giving rise to a functionalist approach where thinking about frameworks, hierarchies, calculation, and quantification dominated. This initial approach is characterised as consistent with Mikes’ (2010) “strategic controller” archetype of the CRO role.
- The challenges to that approach which arose from encounters with a world of engineering and management practice where, perhaps surprisingly, opportunities to accurately represent risk with quantitative precision were relatively rare. In this context risk was predominantly something subjective, ambiguous, and uncertain;
- Which precipitated a shift to a new way of thinking about the role framed by an understanding of Risk Management as a “polarising lens” for reinterpreting the organisation in risk terms. This gave rise to a process-based approach where thinking about knowledge production, knowledge transfer, and facilitation dominated. This latter approach is characterised as consistent with Mikes’ (2010) “strategic advisor” archetype of the CRO role.

Chapter 2 explains how the research evolved out of a particular set of circumstances into a collaborative, transdisciplinary project around this common problem of how to translate generic ERM concepts into value-adding business practices. In this regard, the research did not address a clearly and a priori defined question. Rather the contributions that the thesis makes are the product of work that was undertaken to resolve certain problems and dilemmas, both practical and theoretical, encountered during the course of the project. Those problems emerged somewhat messily out of the initial period of empirical observation

---

<sup>1</sup> Prior to the restructuring of Auckland governance in 2010 Watercare was the provider of bulk water supply and wastewater collection and treatment services to the Auckland region of New Zealand (population approx 1.4 million). Key statistics about the company and a detailed description of Watercare’s corporate context, asset management processes, and risk management framework at the time the research was conducted are provided in Appendix I.



and collaboration, and the period of analysis and synthesis which followed. In this sense, the thesis weaves together various threads around the overarching problem of the CRM's strategic approach to his role.

This chapter locates the thesis within the multi-disciplinary international discourse around ERM and the CRO role in order to frame up the general problem to which the thesis contributes; a task complicated by the fact that there currently exists no coherent overarching framework for the role. The first section below discusses the emergence of ERM and the CRO role and the drivers behind their uptake by industry. This reveals two primary claims about what ERM can deliver in organisations:

- i) ERM supports good governance and internal control (the *governance* claim): ERM serves to align and constrain the interests of organisational actors (staff and management) with those of the principal (the shareholders), and provides assurance of organisational processes through procedural and administrative technologies;
- ii) ERM supports good decision making (the *value-creation* claim): ERM achieves improved enterprise performance through better decision making.

The latter of these claims is a significant theme throughout the thesis. It is argued that public and private sector interest in the benefits of ERM remains high due corporate perceptions of an increasingly risky world and in the wake of the recent global financial crisis.

A feature of the research is the integration of academic and non-academic knowledge to grasp the scope and complexity of a real-world problem of translating ERM concepts into practice. An important contextual point is that, while highly educated, the CRM was a novice CRO. His performance was significant as an experiment with a certain approach, which was ultimately problematic, and for his subsequent reconceptualisation of the role. In order to understand that performance and its significance, it is necessary to elaborate the conceptual world in which the CRM was operating. The second section of the chapter therefore positions the CRO as an emerging professional role and reviews what is known about the role from both academic *and* practice literatures. It is argued that despite the vastness of the Risk Management field and the volume of discourse around ERM, the knowledge base support specific to the task of implementing ERM, and hence specific to performing the role of CRO, as distinct from the task of assessing and managing risk, is very limited. There is currently a scarcity of practical ("how to") and academic (theoretical) guidance for the implementation of ERM and the CRO role. This point both contextualises the CRM's performance and constitutes the general problem to which the thesis contributes.

The third section of the chapter constructs an orienting framework, drawing on the

earlier literature review, which conceptually relates the interventions of the CRO to the effects of ERM implementation in organisations. This framework serves as a heuristic device for demarcating where and how the thesis contributes to the CRO knowledge base. The reader is referred to Appendix III which contains supporting material for this chapter.

## **Origins and drivers of ERM and the CRO role**

Both Enterprise Risk Management and the role of the Chief Risk Officer are emergent features of the global “explosion” of the risk management industry, which began in the early 1990s initially in the financial services and insurance industries (Dickinson 2001; Power 2004, 2007). That “explosion”, initiated by regulatory and shareholder concerns over the governance of banking institutions in the late 1980s and the derivatives disasters of the early 1990s, was marked by a proliferation of regulatory, academic, and practitioner discourse around risk and risk management, particularly on the Internet, and by considerable corporate and public sector interest in what holistic risk management could achieve (Power 2004; 2007; see the review in Appendix III). Faced with institutional pressures to demonstrate improved governance and internal control, the financial industry saw a rapid evolution of ideas, processes, practices, and tools under the umbrella of integrated, holistic, enterprise-wide risk management (Colquitt, Hoyt, and Lee 1999; Conley 2000; Dickinson 2001; Lam 2000; Power 2005b). Indeed, by the turn of the century it was possible to talk about the scale of innovation in the industry as a “revolution in corporate risk management” (Culp 2002).

During this period, corporates in a range of industries began to create the CRO position as a catch-all role responsible for the oversight and management of the new risk management approaches (Burlando 1990; Lam 2000, 2003; Wood 2002). The extent of corporate implementation of ERM and adoption of the CRO role has been the subject of a plethora of annual and biannual industry surveys conducted throughout the 2000s, mainly by multi-national accountancy and management firms such as PWC, Ernst & Young, and KPMG. While the various surveys differ substantially in their scope and focus, making cross-survey comparisons difficult, the results indicate that both ERM and the role of CRO have been widely adopted in a range of industries, although penetration in the insurance and financial services industries is higher than elsewhere (see the review in Appendix III). Of particular relevance, firms in Australia and New Zealand are considered to be ahead of the curve in their commitment to and implementation of ERM (Ernst & Young 2006b, Australia/New Zealand Supplement, p ii; Ward 2006, p 3).

## Drivers of ERM uptake

Two different but convergent pressures may be discerned behind the emergence and uptake of enterprise risk management (Dickinson 2001; Power 2005a; 2007; see the review in Appendix III).

One is responsibility-based, rooted in the corporate governance revolution mentioned above (Power 2005a, 2007). This motivation posits risk management as a technology for the good governance and internal control of the organisation (Julien and Rieger 2003; Lam 2003; Miccolis and Shah 2000; Ward 2006). Briefly, the stream of corporate scandals and failures, which gained prominence in the 1990s and which continue to destroy shareholder value today, have served as a backdrop and motivation for significant shifts in governance values internationally. The resulting revolution in corporate governance has seen the integration of ideas about internal control and risk management as synonymous with good governance and accountability. In this context Enterprise Risk Management has emerged as a codified solution to the problem of compliance with a broad range of stakeholder demands for good corporate governance (COSO 2004; Power 2007; Tarantino 2006).

The other pressure is value-based, where risk management is seen not just as a compliance function protecting shareholder value, but also as a positive force for creating it (Barton, Shenkir, and Walker 2002; KPMG International 2007; Lam 2003; Meulbroek 2002; Nocco and Stulz 2006; see the review in Appendix III). Briefly, corporates are increasingly looking at ERM as more than just a codified solution to the problem of assurance and regulatory compliance. The expectation is that ERM provides the framework to shift the focus of risk and controls from merely preserving value to creating it. While there is a clear internal driver to extract creative value from compliance spending, the overarching logic of the value proposition of ERM is that it leads to improved firm performance through better decision making. Whether this link is real or merely perceived, it has become embedded as an institutional logic at the world level, reinforced and legitimised by the perceptions and actions of powerful actors (e.g. by credit rating agencies such as Standard and Poor's). In this view, risk management has been transformed "from a specialist control side-show to a (shareholder) value enhancing activity" (Power 2005a, p 262).

It is conceptually easy to understand these two pressures as convergent, since practically one implies the other (Power 2005a, 2007). On the one hand, shareholder value created through innovative practices can be rapidly destroyed in the absence of good governance and control, as evidenced by the well known failures and scandals of the 1990s and early 2000s (e.g. Barings Bank, Worldcom, Enron), the more recent rogue trading scandals at Societe

Generale and Credit Suisse involving combined losses of more than £5 billion, and, of course, the recent global financial crisis which has been blamed on inadequacies in risk management (The Economist 2009). On the other, compliance with increasingly onerous regulation is costly and there is pressure to demonstrate the value proposition of advanced risk management systems and practices (Barton et al. 2002; KPMG International 2007; Lam 2006; Marshall, Isaac, and Ryan 2006).

## **Continuing interest in ERM and the CRO role**

Industry surveys published before the global financial crisis of 2008-2009 reported a number of trends which were indicative of a continuing corporate interest in ERM (these are summarised in Appendix III). Following the global financial meltdown of 2008-2009, recent commentary has drawn attention to failures in risk management, and particularly the pervasive over-reliance on little understood quantitative financial and economic models, as an important root cause of the crisis (Taleb, Goldstein, and Spitznagel 2009; Stulz 2009; Champion et al. 2009; Kiviat 2008; The Economist 2009). The crisis has subsequently generated renewed regulatory and corporate interest in risk management, corporate governance, and internal control, and the promotion of ERM is correspondingly at an all-time high (AON Corporation 2009; Beasley, Branson, and Hancock 2009; Coffin 2009a; Dixon 2009; Ignatius 2009). The fall out from the crisis includes, for instance, the December 2009 approval by the American Securities and Exchange Commission (SEC) of new rules requiring public companies to disclose, among other things, their Board's leadership structure and the extent of their Board's role in risk oversight (SEC 2009). Such regulatory tightening creates corporate governance problems to which ERM is hailed as the essential solution (Beasley et al. 2009; Whitcomb 2010). Post-crisis surveys have reported that the current economic volatility and increased attention from regulators are driving a corporate focus on improving ERM capabilities; particularly the alignment of risk management with business strategy and objectives, capabilities to address risk on an enterprise-wide basis, methods for anticipating and representing risk, and the efficiency of risk management departments (AON Corporation 2009; Ernst & Young 2009; Marsh and RIMS 2009). Thus, renewed shareholder and regulatory attention to corporate governance, internal control, and risk management in the wake of the global financial crisis looks set to sustain corporate interest in ERM for the foreseeable future.

Similarly, while there has been the occasional dissenting view that the role of CRO might be a short-lived fad (Ciccarelli 2003; Conference Board of Canada 2001; Fogg 2006; Quinn 2004), a position that no one in their right minds would want to step into (Tuohy 2006), or a

position that becomes part-time or even largely unnecessary once the ERM programme is successfully up and running (Aabo, Fraser, and Simkins 2005), there appears to be growing consensus that the CRO is an important corporate role that is “here to stay” (Atkinson 2007; Economist Intelligence Unit 2005; Jackson 2007; Julien and Rieger 2003; Lam 2001; Mikes and Townsend 2007; Wheeler 2009).

## **The state of the CRO profession**

To summarise Abbott (1988), a formal profession implies a jurisdictional claim over a defined set of practices and an abstracted body of knowledge that professionals draw on to perform those practices. The staking of this jurisdictional claim is an ongoing, constitutive process performed through actual professional work, and through the work of formal organisations which emerge to represent, promote, and protect the profession. The professions (e.g. law, medicine, engineering) make up an interactive system where they compete with each other for jurisdiction over tasks and knowledge. In this ecological analogy, the movements of one profession affect the others, and the evolution of the professions is understood to result from the interrelations between different occupational groups, and, in particular, the ways occupational groups control knowledge and skill. Drawing on Abbott (1988), Power (2005b, 2007) has argued that the position of CRO is an emerging professional role which is engaged in jurisdictional competition with at least three existing professional groups within the organisational milieu:

- Internal regulatory officers, considered collectively. Health and safety officers, and compliance officers, for example, have been around for much longer than the role of CRO (Beaumont, Leopold, and Coyle 1982; Weait 1993).
- Existing “risk professionals”, who already work in risk analysis and management, and who claim a professional capacity to this work. The role of CRO represents an additional category within this group, rather than a replacement of existing roles, but to the extent that the role seeks to co-ordinate and integrate risk management across the enterprise this implies a claim to authority over certain aspects of existing risk management practices.
- General Managers (especially with existing Chief Officers of an organisation, i.e. Chief Executive, Administrative, Financial, and Information Officers). Not only are CROs managers themselves, but their purview encompasses, in principle, every domain of

activity within the organisation. As an internal control agent the role of CRO is to a large extent about influencing the behaviour of others, which immediately suggests the potential for conflicts of authority and legitimacy between the CRO and other managers over how things should be done.

To the extent that the role of CRO does or may seek to reinterpret the practices of these existing roles under the umbrella of ERM, or to establish authority over them (at least in certain respects), this amounts to constitutive, definitional work for the profession of CRO, and implies, both practically and analytically, the demarcation of boundaries and relationships between the roles (Power 2005b, 2007). Abbott (1988) suggested that there are two principal ways in which professions accomplish this: through the control of technique or craft, and through the control of abstract knowledge. Abbott argued that it is the latter characteristic of abstraction which best identifies professions and sets interprofessional competition apart from competition among occupations in general: “...only a knowledge system governed by abstractions can redefine its problems and tasks, defend them from interlopers, and seize new problems.... Abstraction enables survival in the competitive system of professions” (Abbott 1988, p 9). Thus, in Abbott’s (1988) terms, the status of CROs as a bonafide professional group may be judged by three criteria: (1) the degree to which formal representation of the group has been established in the form of national and international associations; (2) the degree to which the profession has demarcated and defended claims to a particular set of practices; and (3) the degree to which the profession has established control of an abstracted system of knowledge as the basis for understanding and developing its professional practices. The following sections review the state of the CRO profession against these criteria.

## **The state of formal representation of CROs**

Although a multitude of institutions have emerged to stake claim to and support risk managers as a professional group in general, particularly in the financial services and insurance industries, Chief Risk Officers do not yet have their own dedicated independent professional representation (see the review in Appendix III). Rather, representation for CROs continues to be provided by the broad range of associations representing risk management and compliance professionals more generally. It is also worth noting that other professional groups have claimed responsibility for implementing ERM, including internal auditors, represented by the Institution of Internal Auditors (Mikes 2010), and management accountants, represented by the Institute of Management Accountants (Shenkir and Walker

2006, 2007). Since, as Mikes (2010, p 74) noted “CROs come from many walks of life, including internal audit, external audit, financial management, business management, and consulting”, there is considerable latitude for a variety of existing professions to lay claim to the CRO role.

## **The state of knowledge base support for the CRO role**

Mikes made the following comment about the professional development of the CRO role: “The ideas and practices of risk management, unlike those of long-established professions, have not yet been codified into a unified domain, leaving chief risk officers with a fuzzy role in corporate governance” (2010, p 81). Both parts of Mikes’ statement are accurate in the sense that risk management does not exist as a unified discipline, and there is, as yet, no unified view on what CROs do, how they do it, and what they need to know in order to do it. However, Mikes’ statement is also a broad generalisation which glosses over a considerable diversity of available knowledge on risk, risk management, ERM, and the CRO role. In order to properly frame up the contribution of this thesis it is necessary to be more precise about the where the gaps in the CRO body of knowledge lie (see Figure 1.1 on page 11).

The study of risk and the practice of risk management have evolved concurrently in a number of different disciplines in the physical, social, and applied sciences (Bernstein 1996; Covello and Mumpower 1985; Recharad 1999; Renn 1998), such that it is now possible to talk about a taxonomy of perspectives on risk (Althaus 2005; Renn 1992). This idiosyncratic development, which means that risk has no single disciplinary tradition of its own (Beck 2004; Corbett 2004; Power 2005b), and the long history of risk as an object of both academic and management attention has produced a multi-disciplinary discourse which is simply immense in terms of its scope and volume. But, for the same reasons, the general body of knowledge on risk and risk management is also very well developed exhibiting a broad diversity of theories and methodologies, a significant breadth of scope and depth of detail, high degrees of coherency and consistency of theories and methods (at least within disciplines), and a high degree of empirical validation. This is represented by the bottom part of the shaded bar on the right hand side of Figure 1.1.

This vast knowledge base is more or less relevant to the role of CRO in a contextual sense. That is, organisations face a myriad of different types of risks for which various disciplinary perspectives and methodologies may be necessary for analysis and management, ERM is promoted and adopted in a wide variety of industries, and ERM is, by definition, cross-functional in perspective, seeking to reveal and evaluate relationships between different types of risk. In the organisational context, the extant multi-disciplinary knowledge base around

risk and risk management therefore informs the analysis and management of risk at the business unit level (i.e. within functional silos; illustrated by the lower box in Figure 1.1).

The important distinction, however, is that CROs are tasked with operationalising ERM and not with managing certain types of risk within the firm. The features which distinguish ERM from siloed risk management (see Appendix V) are those of integration (taking a cross-functional, enterprise-wide perspective of risk, and embedding risk management throughout the organisation, its processes, and its culture), and strategic focus (the explicit centering of the organisation's strategic objectives as the focal point for all encounters with risk). In this regard, CROs are variously described as co-ordinators, advisors, strategists, analysts, synthesists, catalysts for change, developers of best practice, designers and communicators, but *not* implementers (Power 2005b, p 141). Chief Risk Officers commonly have responsibility for setting risk management policy, and may even have significant power in risk allocation decisions, but the responsibility for owning and managing risk ultimately lies with the Chief Executive, and, by delegated authority, with executive and line management of the organisation; and this responsibility cannot be abdicated to the CRO. In this sense, the management of risk in a firm can be seen as divided into three functions, with the latter two fulfilling advisory and support roles to the first: (1) the taking of risk, undertaken by firm management, headed by the CEO; (2) the observation and support of risk management, undertaken by the office of the CRO; and (3) the monitoring and audit of risk management, undertaken by the office of the Internal Auditor (Acharyya 2008). The CRO is therefore not an independent manager of risk who seeks to relieve management of their responsibilities, but is, rather, responsible for ensuring the quality, comprehensiveness, and transparency of the processes and practices in which they engage in order to manage risk (Petit 2006; Power 2005b, 2007; Roberts 2006).

Enterprise Risk Management therefore involves a different set of practices and knowledges than the siloed management of risk, as does the role of CRO when compared to traditional risk managers (illustrated by the upper two boxes in Figure 1.1). Unfortunately, the available knowledge base on the implementation of ERM and the CRO role is substantially less developed than for risk and risk management in general, characterised by relatively little theoretical or methodological development, a significant breadth of scope but little depth of detail, and practically no empirical validation of longstanding assumptions (represented by the shaded bar to the right of Figure 1.1).



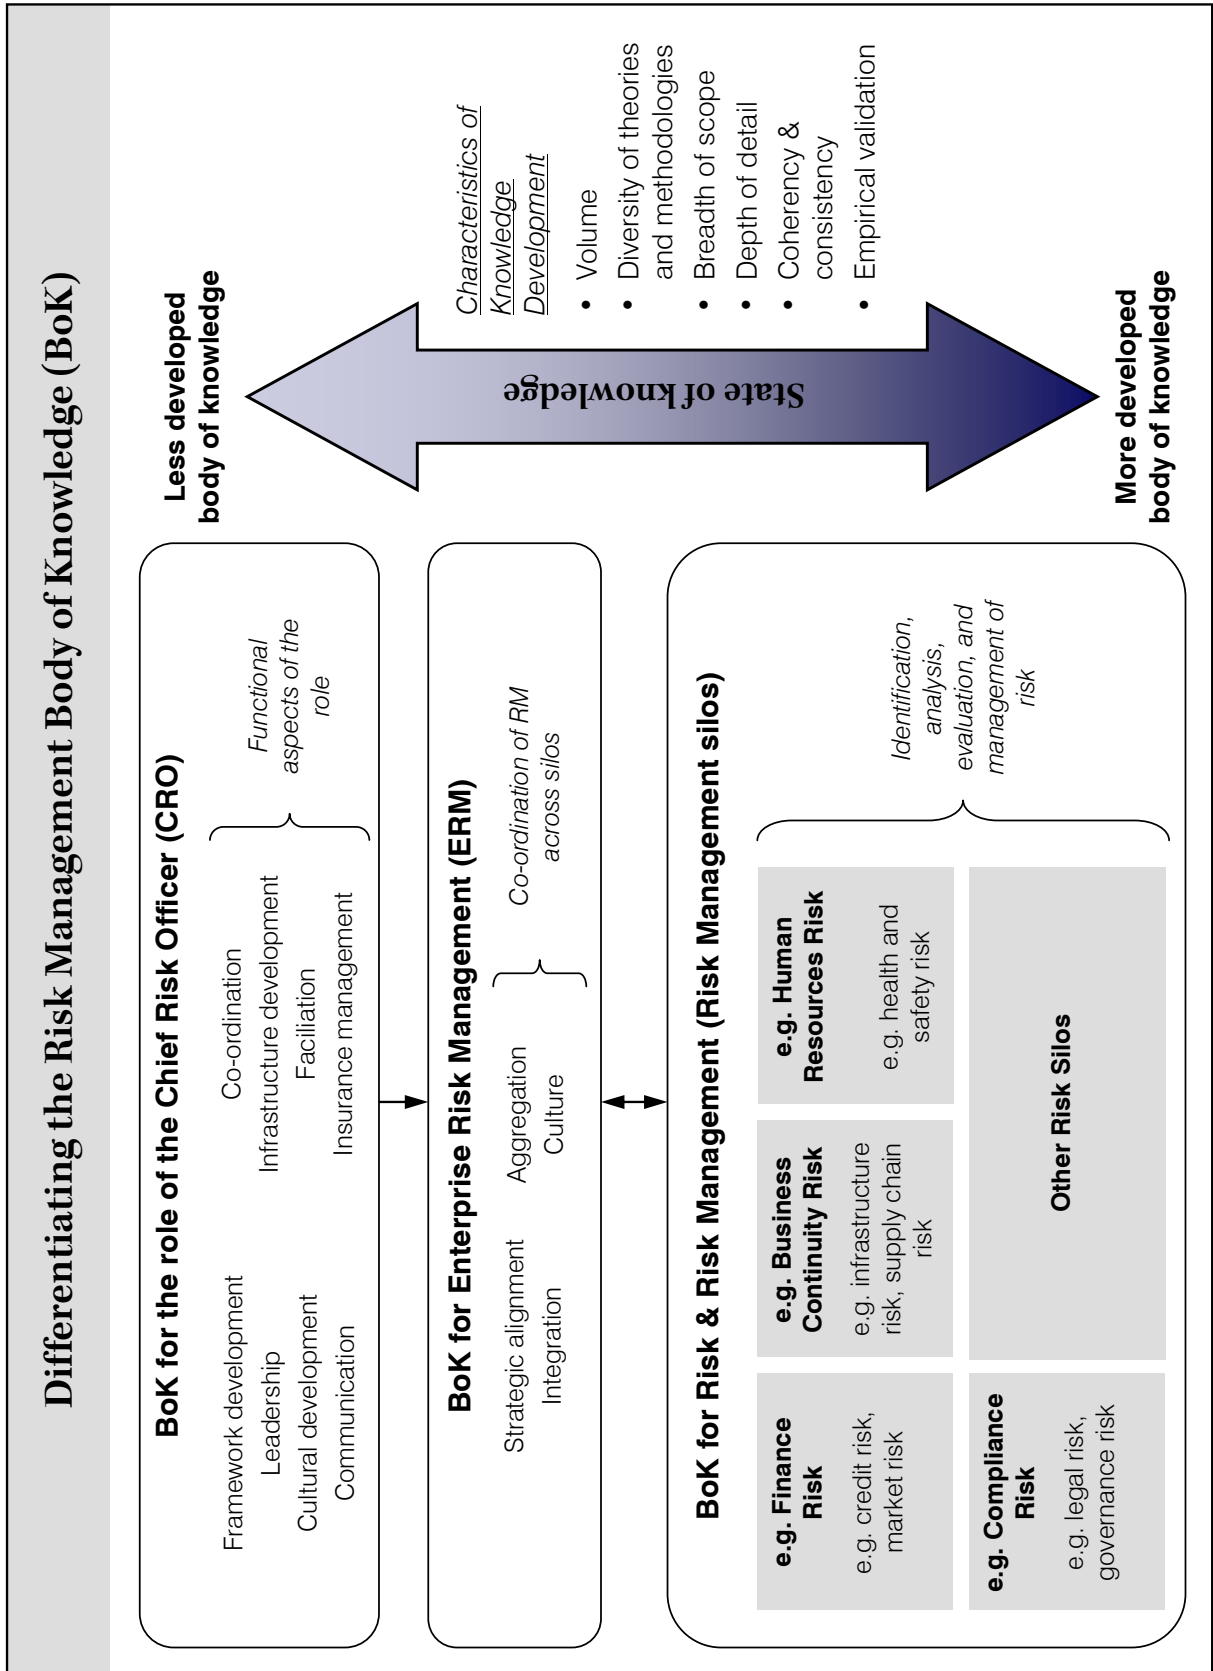


Figure 1.1. Current state of the risk management knowledge base

## Guidance from the practical literature

There is a substantial trade and practitioner literature dealing with the implementation of ERM and the role of the CRO, particularly from the financial services and insurance industries. But, while broad in scope and prolific in terms of volume, this practical literature remains quite shallow in terms of the level of detail and insight provided. It provides extensive coverage of *what* CROs should focus their attention on, such as the importance of defining risk terminology and policies and articulating the benefits of ERM, on the importance of linking risk management to strategic objectives and of cultivating a risk aware culture (see Table III.5 in Appendix III, p 426). And it provides insight into the general nature of the work in which CROs engage, and the attributes and skills necessary to be a CRO. For instance, the nature of the role requires someone with strong interpersonal skills who can build relationships with actors from the Board down to the worker's on the shop floor, and across the breadth of the organisation's functions (see Table III.6 in Appendix III, p 427). Since the CRO is primarily an integrator and co-ordinator, the role requires someone with a breadth of knowledge and experience across risk domains, and who is comfortable working in a multi-disciplinary environment. The CRO must be able to bring together people and information from a variety of sources, facilitate collaborative knowledge production, and then synthesise and communicate that knowledge to others. The ability to integrate information from disparate sources, and to step back and 'see' the strategic implications at the enterprise level (i.e. the big picture) is regarded as essential.

But the practical literature provides relatively little specific guidance with respect to *how* CROs should perform their roles. For instance, at the risk of over-generalising, publications on the implementation of ERM tend to combine basic coverage of the standardised risk management process with generic descriptions of common or particular (e.g. so called best practice) methodologies, methods, and tools for risk identification, assessment, and control, along with extracts from the broader management body of knowledge on topics such as decision making, change management, culture, politics and power, and relationship management. This was something about which Watercare's Corporate Risk Manager commented in our very first Dialogue in June 2007. He commented that the available literature on ERM, and particularly the ERM standards such as COSO and AS/NZS 4360, did not provide much useful guidance for the task of implementation:

...I'd say that the texts are too high level, too simplified. And maybe that's a problem. Maybe I'm trying to work at a level that's tackling too much. I don't know, but definitely the publications that are out there, like the guidance standards and so on, are so simple. They provide a very generic framework, but

where you try to apply that generic framework to a real business, in a way that you are actually trying to get value out of that process you realise that it's not nearly as simple as it's represented. Ideally, you need to take the generic framework and mould it to suit your business. But that requires an understanding of your business, it requires you to analyse your business, to dissect it in some way... [and to understand] what's fundamentally important to the business in the first place. (Dialogue 1, p 245)

The CRM's comments were characteristic of practitioner perceptions more broadly. A 2007 Conference Board of Canada survey of practitioner attitudes toward ERM literature, found that, surprisingly, the COSO framework document (2004) was not used as the key source of information and guidance on ERM, that "...much more work is needed in the areas of research and case studies so that risk executives can learn from the experiences of others who have successfully implemented ERM", and that "...many areas still need to be explored and discussed before a common understanding or methodology for ERM could be considered to be in place..." (Fraser, Schoening-Thiessen, and Simkins 2008, pp 84-85). Responses to the question "What problems/challenges have you encountered in implementing ERM that were not addressed in the literature?" revealed three critical areas of need: (i) how to deal with the myriad cultural, logistical, historical challenges that exist and are unique to all organizations; (ii) a lack of detail with respect to how to integrate risk management across the enterprise; and (iii) that the impact of corporate culture on ERM implementation and practices is not well addressed in the literature (Fraser et al. 2008, p 84). Respondents to the survey also commented that while consultants were often useful during the "getting started" phase, and for particular aspects of implementation, they also tended to offer advice that was too generic or limited in perspective (Fraser et al. 2008, p 78).

The critical finding from Fraser et al. (2008) was that there is a lack of practical and detailed guidance on how to translate the generic process models depicted in the various international ERM standards into effective and value-adding practices in business contexts. Risk executives are "...looking for more practical "how to's", sharing of experiences, impacts of different corporate culture, and best practices at the different stages of ERM implementation" (Fraser et al. 2008, p 84). A recent contribution from Fraser and Simkins (2010), *Enterprise Risk Management: Today's Leading Research and Best Practices for Tomorrow's Executives*, makes an important step forward in this regard, contributing practical guidance on key aspects of the CRO role, with "how to" sections on creating and using corporate risk tolerance, how to plan and run a risk management workshop, how to prepare a risk profile, and how to allocate resources based on risk.

## **Guidance from the academic literature**

Academic literature on ERM is currently very limited, demonstrating a range of idiosyncratic approaches and inconsistent results (Iyer, Rogers, and Simkins 2010). What little research currently exists has so far lacked a foundational framework and appears to have no particular disciplinary home, rather bridging “several business disciplines, including accounting, finance, insurance, and perhaps management and operations management” (Iyer et al. 2010, p 437).

The majority of academic surveys have focussed on the question of why firms adopted ERM, and the relationship between firm characteristics and the extent of ERM implementation (Beasley, Pagach, and Warr 2008; Beasley et al. 2005a; Beasley, Clune, and Hermanson 2005b; Colquitt et al. 1999; Desender 2007; Kleffner, Lee, and McGannon 2003b, 2003a; Liebenberg and Hoyt 2003; Pagach and Warr 2008a). More recent surveys have moved beyond the question of adoption to examine aspects of whether ERM adds value (Gates, Nicolas, and Walker 2009; Gordon, Loeb, and Tseng 2009; Hoyt and Liebenberg 2009; Pagach and Warr 2008b). All of the above twelve research studies represent attempts to statistically relate the presence of ERM to various firm features or characteristics both within and across industries relying on data acquired from surveys of risk managers. Iyer et al. (2010) noted in their review that these studies have so far been hampered by a lack of well-defined variables or indicators which measure either company-level implementation of ERM or the degree of implementation. None of the above studies sought to directly examine the role of CRO.

While such statistical research can go some way toward answering the questions of why firms adopt ERM, what influences the extent of implementation, and whether ERM adds value, it provides little if any insight for CROs in terms of how to perform their roles. Case studies of ERM implementation are of more direct relevance for this purpose. A small number of such case studies have been published (Aabo et al. 2005; Acharyya 2008; Harrington, Niehaus, and Risko 2002; Mikes 2005; Nocco and Stulz 2006; Stroh 2005) and are reviewed in Appendix III. The principal conclusion from that review is that the existing case studies of ERM implementation provide only limited guidance for CROs in terms of how to perform their roles. Nocco and Stulz (2006) provide a strong, coherent, and theoretically grounded framework for ERM implementation, but it is heavily biased toward the calculation of financial risk metrics. Of the other studies, the one by Mikes (2005) stands out because of the methodological approach (70 in-depth interviews with participants at the two banks) and the level of detail provided by the narrative style, including liberal use of direct quotes from participants. Because Mikes targeted her research at “the coal-face” of ERM practice she was

able to show how the performance of ERM in action is intricately tied up with the personalities, attitudes, and culture of the organisation, the pragmatics of doing business, and the idiosyncrasies of corporate history. Rather than attempting to explain ERM in terms of financial theory, she drew upon concepts from organisation and management, and sociology to make sense of what was going on in the two banks. But while Mikes study thus begins to reveal the complexities of ERM implementation, it still contributes primarily descriptive knowledge about ERM implementation, rather than providing knowledge for action (i.e. “how to” knowledge) for the CRO role. The same can be said of Acharyya’s study (2008).

Academic research specifically on the role of Chief Risk Officer is practically non-existent, but three contributions, reviewed in Appendix III, reveal insights about the activities (Ward 2001) and roles (Mikes 2010) performed by CROs, and about the difficult work in which CROs engage to legitimise their profession (Power 2005b). Although Power discusses the CRO role in a general sense, he makes two prominent points which parallel key themes in this thesis. First, he calls attention to ERM concepts as rational designs for risk management systems which reconceptualise the organisation from a risk management perspective (Power 2005b, 2007). Second, by comparison with other functionally dedicated officer roles in organisations, he highlights that CROs are likely to face significant challenges, both political and to their mandate, in articulating the business case for compliance, in mobilising other actors, and in promoting changes to business policies and practices (Power 2005b, 2007).

Mikes (2010) description of three archetypal CRO roles (“compliance champion”, “strategic controller”, and “strategic advisor”), is also significant for its parallels with the CRM’s paradigm shift as a central theme in this thesis. Mikes (2010) differentiated the three archetypes by the degree of emphasis placed on compliance with regulatory and risk management standards, the extent and sophistication of a firm’s risk modelling, and the attitudes of individual CROs. In particular, she suggests that the distinction (which is not black and white) between the “strategic controller” and “strategic advisor” roles ultimately rests on a philosophical choice over “where to draw the line between what can be reliably measured and modelled and what must be placed in the hands of qualitative judgement” (2010, p 79). But she also notes that the roles require very different capabilities. The “strategic controller” role “calls for building a sophisticated risk-modelling capability which is foundational to risk-based performance measurement”, while the “strategic advisor” role “requires an intimate knowledge of the business and what can go wrong – experience that risk officers can only gain by having lived through many organizational successes, losses, and crises” (Mikes 2010, p 80).

However, while the contributions of Power and Mikes begin to expose, describe, and categorise the CRO role, they do not so much contribute knowledge *for* the role as knowledge *about* it. This is the most significant feature of the ERM and CRO knowledge base. The practical literature is vast, but lacks depth, while what little academic research has been published primarily contributes descriptive knowledge about ERM and the CRO role rather than practically relevant guidance. Fraser and Simkins summarise the situation as follows:

I think the research opportunities for enterprise risk management are endless. Little has been written to date in academia on ERM, despite the vast numbers of people and organizations now attempting it. Much of what has been written outside academia has been by consulting firms with their own agenda and marketing motives. It is still an evolving science and therefore case studies and identifying “best practices” is needed... What is succeeding? What do so many fail? There is still mass confusion (Tower of Bable) where there is not even a semblance of alignment among the disciplines (even those present in this discussion) as to what ERM is. (excerpt from Fraser and Simkins 2010, Ch. 26, Journal of Applied Finance Roundtable discussion on current issues and initiatives in ERM, p 501)

In this regard, parallels may be drawn with Morris et al.'s (2006) assessment of the project management body of knowledge. They argued that the project management body of knowledge is characterised by a plurality of perspectives and a lack of integration between research-based and practitioner-based knowledges, and pointed to a large stream of research issues to be addressed, including defining the appropriate scope for the body of knowledge and what paradigm or world view should underpin it. They concluded that research has a real and important role to play in “providing theoretically grounded, empirically-based evidence of the knowledge – and wider aspects of competence – needed to manage projects successfully” (Morris et al. 2006, p 711). It would seem that the ERM/CRO body of knowledge is currently in much the same state, if not actually substantially less developed (Iyer et al. 2010; see also Fraser and Simkins 2010 Ch. 26).

## **Where and how this thesis contributes**

The previous section framed up the following broad problem: that there is currently a scarcity of practical (“how to”) and academic (theoretical) guidance for the implementation of ERM and the CRO role. In order to be clear about where and how this thesis contributes to reducing that scarcity, I employ two heuristic devices to make sense of the body of knowledge (both actual and potential) around ERM and the CRO role. They are (see Figure 1.2, page 21):

- 1) The concept of a general Theory of Enterprise Risk Management (after Chermack

2005), which can be thought of as theory that seeks to describe and explain the firm-level effects of ERM implementation;

- 2) The concept of a Theory of Action for the role of Chief Risk Officer (after Argyris and Schön 1974, 1978), which can be thought of as theory which seeks to describe and explain (and hence inform) the actions of CROs with respect to fulfilling the objectives of ERM.

The two theories are both related and differentiated by a simple conceptual model of the effects of ERM implementation in an organisation (Figure 1.2). The following sections explain the model in Figure 1.2 and the above devices. The reader should note that Figure 1.2 and the following sections do not propose a precise and definitive mapping of the ERM/CRO body of knowledge. Rather the concepts and relationships developed in the following sections are intended only to serve a sensitising function (Blumer 1969). That is, Figure 1.2 and the following sections represent an attempt to establish a general sense of the lay of the land and to make some key distinctions about “what is relevant” with respect to understanding the role of the CRO and the effects of ERM implementation.

## **A Theory of ERM: explaining the effects of ERM implementation**

The earlier review of the emergence of ERM as both a concept and management function, and its uptake in various industries, identified two primary claims in respect of the value of ERM in practice. The first, referred to here as the *governance* claim, is that ERM constitutes a procedural and administrative technology for the good governance and internal control of the organisation, providing assurance of organisational processes, and ultimately improving the reliability of firm performance. The second, referred to here as the *value-creation* claim, is that ERM leads to improved firm performance through better decision making. Both claims establish an explicit link to organisational performance, usually referred to in terms of shareholder value. In essence, when employed effectively ERM should not only protect shareholder value through more reliable firm performance and loss avoidance (the *governance* claim), but also create shareholder value by improving the efficiency and effectiveness of corporate decision making (the *value-creation* claim).

At first glance, the *governance* and *value-creation* propositions appear to be symmetrical with the concepts of “down-side” (threats) and “up-side” (opportunities) risk (Ward 2005). But the categories are not mutually exclusive:

- **Threats** involve decisions about controls, and those decisions may (i) be more or less

defensible in the way they are made, and (ii) involve opportunities for value creation, i.e. some options may be more or less effective and efficient than other options.

- **Opportunities** involve decisions about pursuing and securing certain benefits, and those decisions may (i) be more or less defensible in the way they are made, and (ii) involve risks, i.e. some options may carry a greater or lesser chance of negative outcomes.

The point is that both up-side and down-side risks imply decisions about what to do, and those decisions may themselves provide opportunities for value creation and/or carry additional down-side risk. On this basis, if ERM delivers any value at all then that value must derive from the process itself, not simply from whether risk management targets up-side or down-side risk. This means that the value propositions in the ERM literature must be seen to constitute claims about the firm-level effects of ERM implementation, effects which presumably emerge out of changes wrought by the construction of ERM infrastructures, the performance of risk management processes, procedures, and practices, and the development of the so-called ‘risk culture’ within the organisation. Those claims make a number of assumptions about the effects of ERM implementation, both at the enterprise level and within the firm (Table 1.2).

*Table 1.2. General assumptions about the effects of ERM implementation  
(in the world-level literature)*

<b>Assumptions about</b>	<b>Governance claim</b>	<b>Value-creation claim</b>
Enterprise-level effects of ERM	That ERM is positively associated with the protection of shareholder value, i.e. “improved” firm reliability.	That ERM is positively associated with the creation of shareholder value, i.e. “improved” firm performance.
Internal effects of ERM	That ERM “improves” governance and internal control processes.	That ERM “improves” decision making processes.
Relating internal effects to enterprise-level effects	That “improved” governance and internal control processes are positively associated with the protection of shareholder value.	That “improved” decision making processes are positively associated with the creation of shareholder value.

Read as propositions, the statements in Table 1.2 constitute, in broad conceptual terms, a Theory of Enterprise Risk Management. I have borrowed the idea of a “Theory of Enterprise Risk Management” from Chermack (2005) who used Dubin’s (1978) theory building research method to construct a general theoretical framework linking “scenario planning”, as



particular kind of intervention in organisations, to firm performance; which he called a Theory of Scenario Planning. In Chermack's (2005) version, for example, the intervention was "scenario planning", which was causally associated with "learning", which was causally associated with "mental models", which was causally associated with "decisions", which was causally associated with "organisational performance". I am not claiming to have developed an equivalent theory for ERM. In the first instance, the task of filling in any such Theory of ERM would be complicated by the normative scale of the ERM concept which encompasses a wide range of structural and cultural interventions (e.g. risk frameworks, risk registers, policies, training, etc.) all of which are designed to alter the behaviour of actors in the organisational milieu. Consequently, ERM encompasses a much broader variety of possible interventions than does "scenario planning", which suggests that there may well be a large number and variety of potential hypotheses relating those interventions to changes in enterprise performance. From a normative standpoint the task of deductively explicating hypotheses for a theory of ERM would be substantial; and this thesis makes no attempt to do so.

Rather, my purpose in framing up the concept of a Theory of ERM is to further differentiate, in a heuristic sense, the knowledge base around ERM and the role of the CRO. Specifically, the idea of a Theory of ERM conceptually demarcates a domain of knowledge and theory concerned with what the firm-level effects of ERM implementation are, and correlating those effects with certain types of ERM interventions within the organisational milieu. This is illustrated on the right hand side of Figure 1.2, where the questions in the top and middle boxes (as indicative examples) would be relevant to the development of any potential Theory of ERM. Much of the extant research on ERM has contributed in this area – particularly around correlating ERM implementation with firm performance.

## **A Theory of Action for the CRO: explaining how to perform the role**

In the preceding section I referred to the broad range of structural and cultural interventions encompassed by the concept of ERM. Since the CRO has responsibility for the organisation and oversight of ERM in an organisation, it may reasonably be assumed that the selection and production of those interventions falls under the CRO's remit. This includes the definition and construction of elements of a physical infrastructure within the organisation (e.g. policies, rules, procedures, tools, etc.), and engagement in a variety of practices, processes, and activities (e.g. workshop facilitation, training, dialogue, etc.), which, although perhaps less tangible than the physical infrastructure of ERM are no less important to its

implementation. In Figure 1.2, the term “implementation” is intended to convey this breadth of activity, referring to both what CROs do and what they create within the organisation.

It seems reasonable to assume that when a CRO intervenes in the organisational milieu, he or she is not doing so in a random or ad hoc fashion, but rather is guided by a broader reason and rationale, whether explicit or implicit. To the extent that such a rationale guides the actions of the CRO, it constitutes a Theory of Action (Argyris and Schön 1974, 1978):

A full schema for a theory of action, then, would be as follows: in situation S, if you want to achieve consequence C, under assumptions a . . . n, do A. . . A theory of action is a theory of deliberate human behaviour which is for the agent a theory of control but which, when attributed to the agent, also serves to explain or predict his behaviour. (*excerpt from Argyris and Schön 1974, p 6*)

A comprehensive and detailed Theory of Action for the role of Chief Risk Officer would, ideally, serve both to inform CROs about what they are supposed to do and how to do it, and to explain the actions of CROs in different contexts. As with the idea of a Theory of ERM, I am not proposing to have developed a Theory of Action for the role of CRO. Rather, the concept of such a Theory of Action conceptually demarcates a domain of knowledge and theory concerned with what CROs do and how they do it. In particular, such a theory would be concerned with how CROs choose which interventions to pursue, and how they produce those interventions in specific contexts. This is illustrated on the right hand side of Figure 1.2, where the questions in the bottom and middle boxes (as indicative examples) would be relevant to the development of any potential Theory of Action for the role of CRO. To date there has been little research in this area, although Mikes (2005, 2010) contributions, reviewed earlier, fall into this category.

## **The relationship between the two theories**

The left hand side of Figure 1.2 relates the actions of the CRO to the aforementioned propositions about the intra-firm and enterprise-level effects of ERM. In this sense, the left hand side of Figure 1.2 represents, in crude conceptual form, the real world in which CROs work. The right hand side of Figure 1.2 represents the corresponding conceptual world in which CROs operate, defined by a series of indicative questions, the answers to which would constitute the two theories just described. The two theories are related by a common object (ERM implementation in the firm), but differentiated by scale: the Theory of ERM would be concerned with the enterprise-level effects of ERM, while the Theory of Action for the role of CRO would be concerned with its implementation and development.

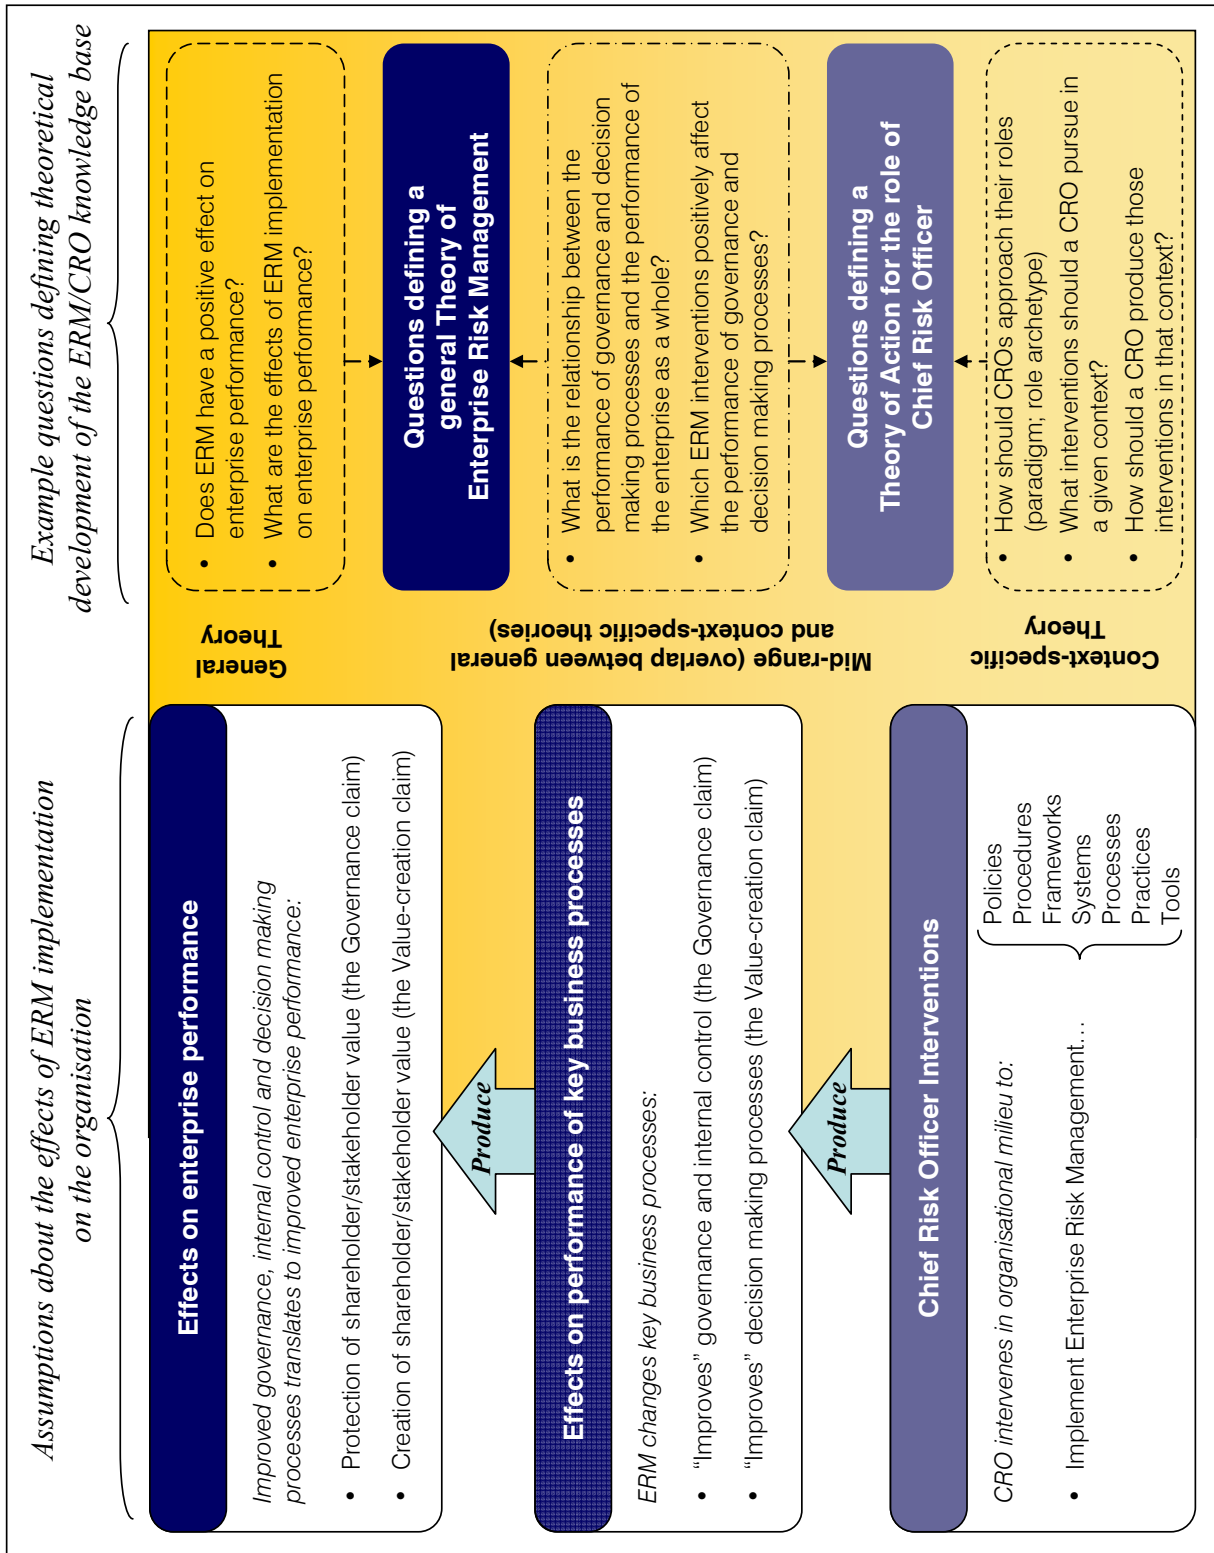


Figure 1.2. Conceptual model relating a Theory of Enterprise Risk Management and a Theory of Action for the role of Chief Risk Officer.

The point is that while a Theory of ERM would clearly be relevant to the role of the CRO, it would not be sufficient as a knowledge base for performing that role because it would lack the “how to” details required to engage in specific organisational contexts to achieve the objectives of ERM. The claims in the ERM literature propose that the positive effects of ERM implementation on organisational performance are attributable to “improvements” in decision making and governance and internal control processes (Table 1.2). This is part of the Theory of ERM, but it provides little guidance to CROs with regard to how to do their jobs, which is the principal complaint expressed by practitioners with respect to the current state of guidance around ERM (Fraser et al. 2008; Fraser and Simkins 2010, Ch 26). A Theory of Action for the role of CRO, in contrast, would inform CROs as to which interventions they should pursue in a given context as being the most effective means of improving decision making or governance and internal control processes in that context. Such a theory would also provide guidance on how the CRO should approach the task of producing those interventions in a methodological sense. The two theories would probably overlap in the mid-range (as illustrated in Figure 1.2).

## **Locating the contributions of this thesis**

The above concepts of a Theory of Enterprise Risk Management and a Theory of Action for the role of Chief Risk Officer serve as heuristic devices which allow me to frame up how the thesis contributes to the problem area outlined earlier. Specifically, the thesis contributes toward a Theory of Action for the role of Chief Risk Officer by theorising an answer to the following principal question: How should Chief Risk Officer’s understand their role with respect to “improving” organisational decision making, and what strategies might they employ toward this goal? The thesis is primarily concerned with what CROs should do, and why, with respect to fulfilling the decision support function of their roles (i.e. that function associated with the value-creation claim of ERM; right-hand side of Table 1.2). The thesis offers the following substantive contributions with respect to this question:

- Insights and understandings about the design of risk frameworks and registers, and their efficacy as instruments for supporting organisational decision making;
- Corroboration of Mikes (2010) finding that the role of CRO may be conceptualised and performed in different ways. The thesis identifies and discusses two perspectives and approaches, which are characteristically similar to those identified and labelled by Mikes (2010) as “strategic controller” and “strategic advisor”. The thesis surfaces

the constitutive assumptions behind the two perspectives.

- A theoretical framework describing a general typology of problem/decision situations, which may be used to interpret organisations as problem solving systems. The framework makes explicit the key general features of the range of organisational contexts in which a CRO might intervene in order to help agents make better decisions.
- Theoretical justification of a generic strategy that CROs may employ to intervene in organisations to support decision making processes, which is grounded in fundamental understandings of decision making behaviour, and which is cognisant of the unique constraints on the CRO position.

The principal claim that the thesis makes is that, at least as far as supporting organisational decision making is concerned, there are different legitimate ways of conceptualising Enterprise Risk Management, of looking at the key objects of a Chief Risk Officer's attention, and thus of interpreting the role of the CRO. Those different ways of "seeing" imply different possibilities for acting, which may be more or less effective with respect to improving the quality of decision making in an organisation. The thesis argues that, of the identified perspectives, one is more likely to be workable as a basis for productive intervention by CROs than the other perspectives. This claim is not made about or in respect of Watercare's CRM, but about the role of CRO in general. It is constituted in two parts.

First, the thesis presents a detailed, longitudinal case study of the implementation of ERM by the Corporate Risk Manager (CRM) at Watercare Services Ltd (Chapters 3, 4, and 5). The case study is significant because of the CRM's principal focus on how ERM could support "better" decision making in the organisation. Empirical grounding is provided by a close following of the CRM's performance, drawing particularly on a series of recorded Dialogues between the CRM, the author, and other participants. The analysis applies the above question as a lens through which to interpret, evaluate, and discuss the CRM's performance. Theoretical grounding is provided by drawing on literatures from a range of disciplines to examine both the particulars of the CRM's performance, and to generalise beyond the empirical narrative. It is revealed how the CRM encountered certain fundamental problems, rooted in his initial conceptualisation of and approach to his role, how he subsequently reconceptualised the management of risk (and hence risk management) as a constitutive feature of the organisation, and how that led to a fundamental reconceptualisation of his role from primarily "strategic controller" to primarily "strategic advisor" (Mikes 2010). The discussion ultimately problematises the CRM's approach, but in doing so also reveals that the

various problems, issues, and dilemmas that he encountered were not just peculiar to his performance, but rather are symptomatic of conceptions of ERM and the CRO role more generally. Second, Chapter 6 organises theory relevant to understanding decision making processes into a coherent systems framework which provides a basis on which to elaborate and evaluate potential strategies for CRO intervention to support decision making in organisations. It is argued that the “strategic advisor” strategy, conceived as a transdisciplinary, facilitative approach, is more likely to be productive than alternative strategies for supporting organisational decision making.

The reader is cautioned not to assume that the contributions of Chapter 6 are the product of a substantively separable inquiry to that presented through Chapters 3, 4, and 5 (i.e. a case study of a Chief Risk Officer in action versus a theoretical discussion on decision making). While Chapter 6 does address certain questions arising out of the case study, methodologically speaking the content of Chapter 6 is inseparable from that of Chapters 3, 4, and 5. In other words, the logic of the presentation of the thesis is not to be confused with the methodology of the underlying inquiry (this is explained in Chapter 2).

The scope of the above contributions is constrained in three important ways. First, the thesis largely ignores the internal control and compliance function of the role (i.e. that associated with the governance claim of ERM; left-hand side of Table 1.2). While this latter function may well be the major component of CRO roles in most organisations, it was simply beyond the time and capacity of this study to theorise two core functions of the role, let alone their interrelation. Second, the thesis does not consider of the question of whose interests the CRO serves. Where the CRO is located in the corporate structure, who the CRO reports to, the nature of the institutional and regulatory environment of the organisation, the internal power structures of the organisation, and the CRO’s own personal and professional values are all factors which will necessarily influence how individual CROs approach their roles. However, the effects of these and other contextual features of the role were excluded from the theoretical scope of the thesis. Third, as regards what constitutes a “good” decision, the thesis considers only the criterion of rationality. There are, of course, other criteria against which the process and outcomes of decision making may be judged (e.g. objectivity, transparency, inclusiveness, fairness, etc.), but it was not possible to consider the implications of such a broad range of criteria within the limitations of this study. The criterion of rationality was given precedence because in the opinion of this author the imperative to “be rational” comes before all others. The thesis does not, however, narrowly apply the assumptions or prescriptions of economic rationality. Rather, I borrow Callon’s (1998a, 1998b, 1999) notion of “framing” as a process of disentangling, from everything else, what is relevant and significant

for the task at hand, as a way of understanding what it means to settle a decision frame under conditions of the rational imperative (i.e. that one should always do what is appropriate with respect to one's beliefs and actions; Rescher 1995).

## Contributions in engineering

The forte of the engineering profession has been and will continue to be the provision of technical solutions to societal problems. In this regard, engineers are held, and hold themselves, to occupy a place of great importance in society (Auyang 2004; Cardwell 1994; de Camp 1977; FIDIC 1990, 2002; IPENZ 2007; NAE 2004; Robbins 2007; UNEP et al. 1997; WFEO 1991, 1997, 2004). However, the profession has also recognised that what engineers do, how they work, and even what they need to know are changing in response to powerful trends in society. In the 21<sup>st</sup> Century, the professional context of engineering is increasingly characterised by five critical factors (Becker 2006; Bronet et al. 2003; Dym 2008; Installé 1996; Jones 2003; NAE 2004, 2005; Schaefer et al. 2008; Spinks, Silburn, and Birchall 2007; Wesner and Dym 2008):

- 1) Explosive technological advancement producing a geometric growth rate of engineering knowledge and resulting in increasing specialisation and the emergence of new cross-over disciplines (e.g. bio-engineering, nano-engineering);
- 2) Increasing customerisation of design, i.e. design driven by and focussed on the particular needs of individual customers and clients, coupled with expectations for rapid turn-arounds;
- 3) The increasingly distributed and multi-cultural nature of engineering work, where the members of project teams may be located across the globe;
- 4) The inter-disciplinary nature of engineering work, bringing together both engineers and non-engineers to address highly complex problems; and
- 5) Increasingly restrictive ecological, economic, political, legal, and cultural constraints within which engineering objectives must be achieved.

In this context, while technical problem solving remains a core part of professional engineering work, pure technical specialisation has become a mass commodity (Wesner and Dym 2008). Emphasis is being placed on the broader professional skills of engineers, i.e. teamwork, communication, outcomes focus, creativity, life-long learning, and awareness of ethical and social responsibilities (Becker 2006; de Graaf and Ravesteijn 2001; Installé 1996;

IPENZ 2009b; NAE 2004, 2005; Wesner and Dym 2008). Consequently, the last 20 years have seen intensifying calls for the reform of engineering curricula, including more emphasis on design theory and skills, the inclusion of more non-technical subjects, and for a new degree structure (see for example Becker 2006; Installé 1996; NAE 2004, 2005). Such calls for reform can be seen, in historical perspective, as another pendulum swing in the long-running debate about the content of engineering education, particularly over the relative weights given to the analytical (science) versus design (art) aspects of engineering in university curricula, and over how much non-technical content engineers should be exposed to in their formal education (Seely 1999, 2005).

In New Zealand, graduate professional engineers (those completing a four-year Bachelor of Engineering degree) are expected to be competent in a specific technical discipline (i.e. must be able to understand and apply “the mathematical and engineering sciences relevant to one or more of the broad engineering disciplines”) and to have foundations in a range of broader professional knowledges and skills (IPENZ 2009a). Of particular note, graduate professional engineers are expected to be able to analyse, formulate solutions to, and manage risks in *complex engineering problems*; being problems which involve wide-ranging or conflicting technical, engineering and other issues, have no obvious solution and require originality in analysis, involve infrequently encountered issues, are outside the problems encompassed by standards and codes of practice for professional engineering, involve diverse groups of stakeholders with widely varying needs, and have significant consequences in a range of contexts (IPENZ 2009a, p 3).

Further, there is some recognition, internationally, that the contrasting trends of (i) the growth in disciplinary specialisation, driven by the geometric accumulation of engineering knowledge, and (ii) the multi-disciplinary complexity of the problems on which engineers work, has created a need for a new type of engineer to continue the tradition of engineering leadership in solving society’s problems (Robbins 2007). Various referred to as “reflexive” (Robbins 2007), “strategic” (Schaefer et al. 2008), or “complete” (de Graaf and Ravesteijn 2001), this future engineer is envisioned as an integrative, transdisciplinary role, whose purpose is not to engage in design in the traditional sense, but rather to facilitate collaboration between engineers and non-engineers (e.g. policy and decision makers, stakeholders, other professions, academia, and the public) and to synthesise innovative design approaches to complex problems.

The new vision for the professional engineer is thus very different from the traditional and now out-dated notion of a technical problem solver. Rather, the Professional Engineer is distinguished from other types of engineer (e.g. Engineering Technologist, Engineering



Technician; IPENZ 2009a) by the competencies necessary to work in, and indeed lead the development of solutions to complex, cross-disciplinary, multi-stakeholder problems. As such, Professional Engineers must have the knowledge, skills, and leadership capabilities to foster productive boundary-spanning dialogue between disciplines (both engineering and non-engineering), and to conceptualise and realise reliable systems methodologies for tackling problems characterised by complexity, uncertainty, and ambiguity across social, technical, and ecological dimensions (Bronet et al. 2003; Installé 1996; Schaefer et al. 2008). Thus, while engineers have always had a need for and been exposed to non-engineering knowledges in their education, the nature of professional engineering work in the 21<sup>st</sup> Century appears to require approaches and knowledges which transcend not only individual disciplines in engineering, but perhaps also the profession of engineering as a whole.

At first glance, neither the subject of this thesis nor its methodology appear to be relevant to the engineering profession. The subject of the study was not an engineering material or method, but rather the performance of an engineer engaged in what appears to be non-engineering work; while the methodology was distinctly atypical of the positivist methods usually employed in engineering research. However, Watercare Services Ltd was an infrastructure management organisation where engineers, engineering work, and engineering ways of thinking predominated. The subject of the research was an engineer moving into a managerial position, and, at least initially, approaching his new role in a manner characteristic of those ways of thinking. And the research was conducted by an engineer, toward a PhD in Civil Engineering. Chapter 7 reflects on how, as a result of this positioning, the thesis contributes to the current debate about the appropriate scope of education for engineers in the 21<sup>st</sup> Century. It is concluded that the thesis offers the following contributions:

- An example of how transdisciplinary research into professional roles can contribute to the development of the professional knowledge base. Similar approaches could be used to study how professional engineers work, as distinct from what they work on or with, and thus to orient and develop the formal curriculum for the broader education of engineers;
- Transdisciplinary methodology for both the study of engineering roles (Chapter 2), and for performing the kind of cross-disciplinary work that will characterise the “strategic” engineering roles described above (Chapter 6).

## Organisation of the thesis

The functions of chapters 1 and 2 have already been described. The work performed by the other chapters (3 – 7) is summarised below.

Chapter 3 is the first of three chapters which examine the experiences of Watercare’s new Corporate Risk Manager as he sought to implement Enterprise Risk Management within the firm. The purpose of the chapter is to reveal and characterise the CRM’s initial approach to his role, and to justify that approach as reasonable despite his status as a novice Chief Risk Officer. Empirically Chapter 3 lays the groundwork for the analyses in Chapters 4 and 5 by framing up what the CRM wanted to achieve and how he intended to do it (i.e. specific proposals). Analytically, Chapter 3 lays the ground work for the discussions in Chapters 5 and 6 by making explicit (i) the underlying rationale which motivated the CRM’s initial approach, and (ii) the corresponding normative agenda of ERM as a programme for the (re)construction of rational management in organisations.

The CRM’s vision for developing Watercare’s Risk Management capabilities was dominated by proposals for improving the quality of the company’s “risk data”. Chapter 4 describes and analyses the CRM’s experiences as he engaged in certain tasks toward the implementation of that vision: (i) redesigning the company’s risk framework and registers, (ii) promoting quantitative risk modelling, and (iii) facilitating strategic risk assessment. The main points of argument in the chapter concern the descriptions and explanations of what happened in each case, what those experiences reveal about risk as an object of inquiry, and how those insights, in turn, explain the CRM’s experiences. Empirically, Chapter 4 sets the scene for Chapter 5 by revealing how the CRM was relatively unsuccessful in his efforts to improve the quality of Watercare’s “risk data” through the design of ERM infrastructure, as he had initially envisaged, but did achieve success by actively facilitating and guiding inquiry into risk. Analytically, the discussion in Chapter 4 lays the groundwork for that in Chapter 5 by (i) framing up the conventional theoretical distinction between Risk and Uncertainty, (ii) explaining the CRM’s failures as the result of a conflict between expectation and reality over how detailed “risk data” should be, and (iii) explaining the CRM’s success as the product of his alternative facilitative approach.

Chapter 5 presents the discussions between the CRM and myself which centred around two core problematics of the CRM’s experience: (1) What is the value of detailed “risk data” (and hence detailed risk assessments) in contexts of relatively low uncertainty about system performance or decision outcomes?; and (2) How can CROs translate the abstract definition of the Risk Management process depicted in world-level ERM standards into a value-adding

practice in organisations? In each case the CRM reconceptualised Risk Management and ended up redescribing his role in terms which I characterise as a methodological paradigm shift. The discussion section calls the CRM's approach into question by framing up the central dilemma of the CRM's experience. The chapter then draws on theory to argue that the CRM's dilemma was not merely peculiar to his performance, but rather is symptomatic of conceptions of ERM and the CRO role in the broader literature. Chapter 5 establishes the focus for Chapter 6 because it hypothesises about, but ultimately leaves open the question of how CROs should approach the decision support function of their roles.

Chapter 6 addresses from a theoretical perspective the question of how CROs might fulfil the decision support function of ERM. The chapter takes as its starting point the following assumption: if Risk Management is fundamentally about making good decisions, and if the CRO is that corporate agent tasked with organising Risk Management on an enterprise-wide basis, then the strategies, methodologies, and tools that CROs employ in their role should be grounded in a theoretical framework which accounts for how agents make decisions. In other words, if Risk Management is essentially a procedure for helping agents to make better decisions, and the CRO is, in this regard, a kind of Decision Engineer (March 1978), then he or she requires an understanding of how people make decisions and the various factors which influence decision making in practice. The objective of Chapter 6 is to formulate such an understanding, in the form of an orientating framework, as a basis for conceptualising strategies by which Chief Risk Officers may productively intervene within organisations to support decision making.

Chapter 7 reflects on the substantive conclusions of the thesis, the contributions that the thesis makes to the CRO knowledge base, and the directions in which future research might extend and develop those contributions. The chapter also reflects on how the thesis contributes to the future development of the engineering profession.



# Chapter 2

## Methodology

So, judged by the standards of formal method, this book is the product of a messy hermeneutical approach to hunches and intuitions, layered upon each other like strata in a (slow) process of accretion. The method is one of progressively and continuously refashioning a family of arguments in a process of self-critical writing, a constant and restless trade and exchange between theoretical precepts and bits of the empirical world, each mutually explicating the other in a 'to and fro' process as Habermas once put it (*hin- und hergerissen*). This repeated process of 'fitting' theory and matter never ends in principle, but the writing of a book must stop at some stage.

*Excerpt from the preface to Michael Power's book  
Organized Uncertainty: Designing A World of Risk Management (2007, p ix)*

Contrary to traditional and perhaps idealistic notions of scientific method, the methodology of the inquiry was not pre-given. It did not follow some a priori plan to address predefined questions or hypotheses, but rather evolved rather messily out of the circumstances under which I came into the research, the collaborative nature of the initial work of shadowing Watercare's Corporate Risk Manager in 2007 and 2008, and the uncertainty which characterised our respective projects. Consequently the inquiry cannot be located in any single disciplinary tradition.

This chapter frames the inquiry as an experiment in transdisciplinary research. Although my initial approach was an intuitive response aimed at mitigating the risks of my starting point, it evolved into a prolonged collaborative engagement with Watercare's CRM. The chapter discusses this co-productive endeavour in terms of two distinct but complementary projects, convergent around a common interest: the search for understanding of the role of

the Chief Risk Officer. The methodology of the thesis is elaborated by appeal to pragmatist epistemology, which has recently seen a resurgence of interest in the social sciences as a way to sidestep the longstanding philosophical dualisms of disciplinary approaches. The first section, below, presents the pragmatist perspective, focussing particularly on its non-representationalist view of knowledge. It is argued that pragmatic, pluralist approaches offer a way to cope with the uncertainty which attends the crossing of disciplinary boundaries because they are intentionally open to the possibilities of alternative perspectives, and seek to reflexively question and re-articulate preconceptions through dialogical confrontation.

The methodology of the inquiry is reflected in the presentation of the thesis to a certain extent: the organisation of the empirical narrative around the CRM's paradigm shift, the interplay between contextualisation and generalisation, and the dialogic presentation of the theoretical discussion in terms of evolving assumptions about the key objects of attention (this is explained later). It is, however, difficult to get a sense of the nature of the research process from the final document. The excerpt from Power's (2007) book *Organized Uncertainty*, above, provides that sense in that it accurately describes both the experience of and methodological approach to producing the thesis. It expresses the messiness of the experience – the relative lack of formal methods, the playing out of hunches and intuitions guided by the often fleeting perception of meaning between encounters with the empirical and the theoretical, and the slow, accretive process of making sense of the situation. And it conveys the centrality of a tangled and tortuous process of writing to the research method – of slowly tracing back perceptive leaps of intuition, of iteratively explicating and fitting together bits of theory and empirical matter, and of continuously remoulding the arguments until the thesis finally became clear. In this sense, the thesis emerged, not chapter by chapter, but rather as a collection of “bits” which were gradually moulded into a coherent whole (which also serves to explain why no part of the thesis has yet been published).

## **On performing reliable knowledge**

Pragmatism is a broad philosophical movement which emerged in America in the late nineteenth and early twentieth centuries, and is associated with the claim that the truth of a proposition is to be found in the practical consequences of accepting it (Hookway 2010; McDermid 2006). Pragmatism rejects the Cartesian and Lockean ideas that the mind of the knower is a blank slate, upon which the knower, standing aside from the world as an impartial observer, perceives a duplicate copy of the world (Hookway 2010; McDermid 2006).

This “spectator theory of knowledge” conceives of knowledge as the inner representation of the outer world, the truth of which is a matter of the correspondence or fidelity of that representation with reality (Baert 2003). In contrast, pragmatists insist that we should avoid the conceit that human knowledge could ever represent the universe as it actually is (Baert 2003; McDermid 2006). This is more than a belief that the totality of objective reality will always be richer, more complex, and more novel than we can ever grasp through sensory experience or conceptualise within the limits of our linguistic, theoretical, and cognitive resources (Hookway 2010; Rescher 1995, 2005). It is, further, the belief that we cannot ever conclude the correspondence between our knowledge of the universe and the universe itself because we do not, and cannot ever, have completely unadulterated access to that reality (McDermid 2006; Rescher 2005).

This stance flows from the view that perception is very much an active process for the mind. Perception is never simply the raw imprint of sensory information, but is always mediated by distinctions of the mind – we always “see” the world through the lenses of conceptual schema, and never as it actually is (Baert 2003; Hookway 2010; McDermid 2006). Thus knowledge of the world is never simply “given” to the mind through experience, but is always, which is to say even at a fundamental level, the product of inference, of interpretation through conceptual (theoretical) categories<sup>2</sup>: “For rationalists reality is ready-made, complete and waiting to be discovered. For pragmatists it is always in the making” (Baert 2003, p 93; Hookway 2010; McDermid 2006). Knowledge is, in this sense, constituted in a subjective relation between the knower and the known, such that the two cannot really be distinguished (Baert 2003, 2005; Bosch 2007; Johnson and Onwuegbuzie 2004; McDermid 2006).

Pragmatism similarly denies that there is an absolute metaphysical foundation to all knowledge; that all knowledge, in all settings, cultures, and times can be justified by reference to an immutable, all-encompassing framework of categories (Baert 2003, 2005; McDermid 2006). Rather, pragmatism accepts that there is more than one way to encounter the world (e.g. Gardner 1983, 1999), and more than one way to make sense of our experiences (e.g. Pepper 1942), and thus that the schematic possibilities for conceptualising the world and its contents are multiple (Baert 2005; McDermid 2006).

In this regard, pragmatism does away with the idea of an absolute Truth about the

---

<sup>2</sup> Even if it may be the case that there are certain fundamental or foundational distinctions “given” to the mind through sensory experience upon which all other knowledge is constituted (perhaps, for instance, the perceptual distinction between light and dark; Herbst 1976; Spencer Brown 1969), pragmatists would challenge the idea that we can gain access to those distinctions, since any such access would necessarily be mediated by higher knowledge.

universe, or at least the idea that we can gain access to it: on the one hand, we can never conclude the correspondence of propositions with empirical reality because we do not have totally impartial access to it; and, on the other hand there, is no immutable metaphysical foundation on which we can call to justify our beliefs. But this does not lead to hopeless nihilism or relativism. Assuming away the possibility of knowing the absolute Truth of things does not imply that we cannot know those things at all or in reliable ways.

### **There is nothing so practical as a good theory**

In the pragmatist view, knowledge takes the form of “warranted assertions” and “workable lines of action”, understood as the products of inquiry, of problem-solving processes through which the inquirer “struggles” to replace doubt with “settled belief” (Hookway 2010; McDermid 2006; Morgan 2007). Since “people cannot escape using a conceptual system”, there is no pure or original state of absolute naivety or “doubt about everything”, no place to stand which offers the “God’s-eye” view; there is only the “agent’s point of view” (Baert 2003, p 94; McDermid 2006). Thus inquiry always proceeds from somewhere – always begins “in the middle of things” – which is to say that it is always grounded in presupposition (McDermid 2006; Rescher 2005). Indeed, presumption is pragmatically essential to rational inquiry and communication (Rescher 1995, 2005). That is, we can only proceed rationally with our cognitive and practical endeavours *after* we have made certain foundational assumptions about the nature of the world into which we are inquiring; e.g. that there is a reality “out there”, that it is perceptible in an intersubjective sense (“you can see what I see”), and that it has certain characteristics (whether it is a hard, objective reality, independent of what humans might think about it; or multiple, fluid realities constructed through the shared consciousness of human perception) (Rescher 1995, 2005). In the absence of such presuppositions, we would be assuming away any correlation between our knowledge of the world and the world itself, such that we could not reasonably act in the world, nor draw conclusions and learn from our experiences, nor communicate with others in a meaningful way (Rescher 1995, 2005).

The relevant question, then, is how do we do this? Having taken for granted that the world around us is real and that we are observing it, how do we move from the inherent subjectiveness of sensory experience – “I take myself to be seeing a cat on a mat” – to an objective factual claim – “There actually is a cat over there and I am looking at it” (Rescher 2005, pp 28-30)? Rescher’s suggestion (2005, p 30) is that this transition is automatic – in practice we move immediately and unproblematically from observation to fact – and is not



made on the basis of inference (evidence), but rather on the presumptive, pragmatic policy of trusting one's own senses:

This policy itself is based neither on wishful thinking nor on arbitrary decisions: it emerges in the school of praxis from the consideration that a long course of experience has taught us that our senses generally guide us aright – that the indications of visual experience, unlike, say, those of dream experience, generally provide reliable information that can be implemented in practice. (excerpt from Rescher 2005, p 29)

Since we have presupposed reality (i.e. “it is plausible, a priori, that there really could be a cat on a mat”), and we have learnt from long experience to trust our senses, the leap we make from “I see a cat” to “there actually is a cat there” is made on the reasonable presumption that our senses are providing us with objective information, this time, as in the past (Rescher 2005, p 30). But if knowledge is presumptive in this way then what is its truth value? The pragmatist view is that truth is a question of experience, in both historical and future senses.

In the future sense, the pragmatic maxim is that we should put our knowledge to the test by acting on it (Fendt, Kaminska-Labbé, and Sachs 2007). Since we can neither measure the correspondence of a claim with reality nor establish its foundation in some metaphysical paradigm, we can only judge the truth value of a claim by how well it “works” as a basis for action (Baert 2003; Hookway 2010; Johnson and Onwuegbuzie 2004; McDermid 2006; Morgan 2007). Theories “are to be judged primarily by their fruits and consequences, not by their origins or their relations to antecedent data or facts” (McDermid 2006). This maxim that we should attend to the practical consequences of theories has a verificationist character (Hookway 2010), but the pragmatist does not fallaciously interpret the practical consequences of action as evidence verifying or refuting the fidelity of an idea with the ontological nature of its object. Rather, for the pragmatist inquirer, the truth of an idea is expressed in its problem-solving power, as an increase of the power to make the world respond to his or her cognitive interests: “what is more important is whether the idea of theory is successful: that is, whether it accomplishes what one wants to achieve” (Baert 2003, p 95; 2005; Joas 1993; Johnson and Onwuegbuzie 2004; McDermid 2006).

If the future truth value of a knowledge claim lies in its utility or expediency for us going forward, then the historical truth value of a knowledge claim lies in our prior experience with putting it into action (Rescher 2005). Here, “experience” refers not to the immediate sensory experiences of the individual, but rather to long-run “historical experience in its larger transtemporal and transpersonal sense” (Rescher 2005, p 13):

What is essential is that theories pay their way in the long run—that they can be relied upon time and again to solve pressing problems and to clear up significant difficulties confronting inquirers.

To the extent that a theory functions or “works” practically in this way, it makes sense to keep using it—though we must always allow for the possibility that it will eventually have to be replaced by some theory that works even better. (excerpt from McDermid 2006)

Thus theories become true and are true to different degrees according to how well they work, but truth in this sense is always tentative and provisional, open to revision depending on what happens next time (Johnson and Onwuegbuzie 2004, Table 1, p 18). All beliefs and theories “are best treated as working hypotheses which may need to be modified—refined, revised, or rejected—in light of future inquiry and experience” (McDermid 2006). In the pragmatist view the accumulation of knowledge is therefore a lengthy intersubjective process of experiential trial and error where successes and failures serve to direct our attention such that, over time, we come to rely more extensively on certain ideas than others, gradually building up reliable knowledge of the world in which we then place our trust such that we may operate effectively (Constant 1999; Rescher 2005, p 34-35; Rosa 1998).

## **From correspondence to resistance: the objectivity of strong networks**

The pragmatist view of knowledge and truth has a certain affinity with the Actor-Network Theory (ANT) view that any object we may perceive in the world, including knowledge, is an effect produced by a heterogeneous network which as a whole constructs and supports that object (Callon and Law 1997). In ANT, objects are referred to as “actants”, which is a term meaning, literally, something which acts, and which may be used to refer to anything – human or non-human, tangible or intangible, material or conceptual. Although counter-intuitive, the definition is analytical rather than literal. That is, for the purpose of analysis, ANT assumes that people are not necessarily special in comparison to non-humans in terms of their power to act and influence events and other actors (Law 1992; Latour 2005). An actant is simply “something that acts or to which activity is granted by others. It implies no special motivation of human actors... An actant can literally be anything provided it is granted to be the source of an action” (Latour 1996, p 370). Every actant (object, system component) in the world is related (associated, connected) to others, forming what is termed actor-networks, and it is through these relationships (associations, connections) that the effects of action (forces, energy) are transmitted. Some actants act as intermediaries, which merely pass on a force without transformation, while others act as mediators, which transform, multiply, or absorb forces (Latour 2005). Analytically, it is the mediating actants which are of interest since it is they who alter the network and its effects.

From the ANT perspective, all objects in the world are produced, sustained, and ultimately destroyed by, or rather through the networks of relations in which they are embedded (Law 1992, 1999). The terms “performed” and “enacted” are the ANT way of saying that nothing simply exists, that everything is the effect of action: “knowledge, agents, institutions, organizations, and society as a whole, are *effects*, and... such effects are the result of relations enacted through heterogeneous networks of humans and non-humans” (Bosco 2006, p 136). This applies to all objects, whether we perceive them as stable and independent (such as a car), or fleeting and ephemeral (such as ideas and concepts). Thus ANT treats knowledge in the same way that it treats reality – as something which is produced and stabilised through performative processes of network building:

...[knowledge] is the end product of a lot of hard work in which heterogeneous bits and pieces -- test tubes, reagents, organisms, skilled hands, scanning electron microscopes, radiation monitors, other scientists, articles, computer terminals, and all the rest... are juxtaposed into a patterned network which overcomes their resistance. In short, it is a material matter but also a matter of organising and ordering those materials. So this is the actor-network diagnosis of science: that it is a process of "heterogeneous engineering" in which bits and pieces from the social, the technical, the conceptual and the textual are fitted together, and so converted (or "translated") into a set of equally heterogeneous scientific products. (*excerpt from Law 1992, p 2*)

The term “performative” also means that the actions and connections which constitute an actor-network must be continually made and remade, or performed, in order to sustain it. The following distinction is illustrative: Boeing 747s do not fly; airlines fly (Latour 1999). That is, the plane itself is not capable of flight outside the multitude of relations which make flight possible. Without airports, fuel, pilots, ground crews, maintenance and administrative staff, and of course fare paying passengers, the plane is simply an arrangement of materials sitting on the tarmac (but even that particular arrangement of materials, which we may identify as a Boeing 747, is itself an effect of a heterogeneous actor-network, i.e. the network of people, knowledge, organisations, and machines which designed and produced the aircraft). Only when the plane is connected to all those other agents is it possible for the plane to take to the air, but in that case, it is not the plane, alone, that is flying.

Thus ANT refuses to take objects, or, more specifically, the stability and durability of objects in the world for granted. If something is perceived to persist in a stable state, or to resist the actions of others, or to be a source of causality, whether the object in question is knowledge or an idea, an institution, a material artefact or technology, or even a human actor, then ANT suggests that such states of affairs is never simply given, and we should find out how such stability, durability, and causality is achieved (Latour 2005; Law 1992, 1999).

The answer to this question is to be found by tracing the associations between actants, between mediators and intermediaries, to reveal the translations which established those associations, and which sustain the perceived stability, durability, and causality: “If some causality appears to be transported [between two mediators] in a predictable and routine way, then it’s the proof that other mediators have been put in place to render such a displacement smooth and predictable” (Latour 2005, p 108). Translation is the name given to the process by which two (or more) mediators are induced into association, which involves an alignment of interests (Latour 1987, 1988). The term is analytical and does not necessarily imply conscious intent (i.e. as one would commonly impute to human actors). Rather the term calls attention to the fact that in order to get any actant, human or non-human, to do our bidding, to support that which we wish to produce, certain work has to be performed to enrol that actant to our cause (Latour 1987, 1988). In some cases, this translation may be significant – the scientists at CERN have spent vast sums of other people’s money to build the largest and most technically complex machine in the world, all in order to gain control over the world’s smallest elemental particles, such that they may direct those particles to travel in certain directions, at certain velocities, and eventually to collide with each other (this being entirely contrary to the natural “interests” of those particles).

This should not be taken to imply, however, that anything goes, that we can simply “perform” into being whatever we want. For instance, just because we have a conceivable (logically possible) theory of the world, there is no guarantee that we will ever be able to make the world respond to it:

If what I have said is right, nature undoubtedly responds to the theoretical predispositions with which she is approached by the measuring scientist. But that is not to say either that nature will respond to any theory at all or that she will ever respond very much. Re-examine, for a historically typical example, the relationship between the caloric and dynamical theory of heat. In their abstract structures and in the conceptual entities they presuppose, these two theories are quite different and, in fact, incompatible. But, during the years when the two vied for the allegiance of the scientific community, the theoretical predictions that could be derived from them were very nearly the same... It follows that any measurement which... "fit" one of these theories must have "very nearly fit" the other, and it is only within the experimental spread covered by the phrase "very nearly" that nature proved able to respond to the theoretical predisposition of the measurer. (excerpt from Kuhn 1961, pp 176 - 177)

Kuhn’s point in this passage is that if knowledge is constituted in the response of the world to the questions that we choose to ask of it, then truth (i.e. the degree of “fit” between theory and reality) is necessarily conditional on the context and resolution with which we engage the world. It depends on our ability to put the questions to Nature, to “push Nature

around” in ways according to theoretical preconception, *and* on the precision with which we can record (i.e. measure) Nature’s responses: “Only under these conditions and within these limits can one expect nature to respond to preconception” (Kuhn 1961, p 177; McDermid 2006). In other words, knowledge and truth are very much practical matters and, precisely because of this, cannot be arbitrarily or capriciously performed into being (Latour 1987, 1988, 2005; Law and Singleton 2000; Law and Urry 2004). To the contrary, if our claims are to be taken seriously, and relied upon as a basis for action, then they will not be easy to make:

...the process of building a network that will create a sense of reliable knowledge of a real world is also a performance – and one that... cannot be created out of nothing. Raw materials have to be put in place – and then held in place... ‘allies’ (people, facts about the world, laboratories, scientific papers, publishers, instruments, scientific funding agencies, colleagues, referees) have to be cajoled, seduced, bought or forced to play the roles allocated to them. And it is a little worse than this, because it isn’t enough to pick this list of potential allies off one by one. To pick any of them off you have to have most of the others already lined up. Which is a way of saying that they all have to perform together – and if they don’t, if one bunch of actors goes off script, then the network holding all the others in place is also disrupted, and they too are in danger of going native... The argument, then, is that performances are difficult to put on unless they build on the networks that are already in place. That realities and knowledges cannot capriciously be performed into being. That we are, in general, somewhat stuck with what passes for the world, and our knowledge of the world. Which in effect, though not in analysis, produces results that are consistent both with the realist sense that there is a world and that we approximate towards knowing it well, and the pragmatist intuition that knowledges change as we approach the world with different questions in mind. (excerpt from Law and Singleton 2000, pp 3-4)

In this regard, the ANT view of knowledge shares a number of points in common with the pragmatist view outlined above; including the idea that knowledge is never simply given to the passive, receptive mind, but is rather constructed through the intentional, directed action of inquiry; and that success (truth) is a practical matter, expressed in the power of the inquirer to make the world respond to his or her cognitive interests; and that the difference between reliable, objective, or true knowledge, and claims which are subjective or false, is not one of an ultimately immeasurable correspondence to reality, but rather of the strength and durability of the networks which constitute and support them (Callon and Latour 1981; Latour 1987, 1988; Law 1992). Take, for example, the now well established physical law that for a fixed mass of ideal gas at fixed temperature, the product of pressure and volume is a constant (Boyles Law). In the future sense, the truth of this statement lies in its practical utility; that if we were to subject a gas to certain pressures and temperatures then we could expect it to respond in certain ways. In the historical sense, the fact that today we take Boyle’s Law as a universal truth (a “leviathan”) is entirely the effect of a vast network of

associations which first established the relationship as scientific ‘fact’, and then made that fact portable so that it could be embodied in all the tools, instruments, techniques, and machines that we now take for granted (Callon and Latour 1981; Kovach 2004; Latour 1987, 1988). Boyle’s Law may be a “relational effect”, but it is also real and obdurate because it was and still is “produced in thoroughly non-arbitrary ways, in dense and extended sets of relations” (Law and Urry 2004, pp 395-396).

## **The “pragmatic turn” in the social sciences**

Paradigms are particular world views which embody a set of assumptions or beliefs about the nature of the world, about how we may come to know that world, and about which questions are worth asking in a research field, and which procedures are most appropriate for answering those questions (Creswell 1998; Jackson 2000; Kuhn 1962; Morgan 2007). Disciplinary inquiry is always framed within the boundaries of a relevant paradigm; which is to say that the paradigm is prior to the research questions and methods, whether by explicit design or implicit presumption or both (Feilzer 2010; Johnson and Onwuegbuzie 2004). Paradigms resolve for the researcher which of the many legitimate ways of encountering and knowing the world should be given precedence, providing a priori ontological, epistemological, and methodological positions from which to approach and produce knowledge of particular aspects of the world.

In the social sciences there is a broad range of paradigmatic perspectives to choose from (Burrell and Morgan 1979; Jackson 2000). A prominent polarisation is discernable between those who approach the social world in much the same way that physical scientists approach the natural world, exhibiting the assumptions and methods of the realist/positivist paradigm; and those who assume that the social world has a much more fluid and contingent existence, exhibiting the assumptions and methods of the constructivist/interpretivist paradigm (Jackson 2000; Johnson and Onwuegbuzie 2004). These two paradigms are often referred to as Quantitative and Qualitative respectively, referring to the nature of the research methods which are commonly associated with each (Johnson and Onwuegbuzie 2004; Morgan 2007); although others have noted that research methods are not inherently tied to any one philosophical orientation and the production of knowledge in any paradigm may involve qualitative or quantitative work, or a mixture of both (Hatchuel 2000; Kuhn 1961). A second paradigmatic polarisation is discernable, defined according to whether priority is given to understanding the social world as a means for regulation and control (e.g. functionalist, interpretivist approaches) or to open up possibilities for change and reformation (e.g.

emancipatory, post-modern approaches) (Jackson 2000). The extreme forms in each dimension are opposite and incommensurable, and difficult to defend as general perspectives; thus, they are usually seen as conceptual ideals defining a spectrum of paradigmatic positions from which to approach the social world (Jackson 2000; Johnson and Onwuegbuzie 2004).

A significant “pragmatic turn” has been taking place over the last two decades, with researchers from across the social sciences appealing to pragmatist epistemology in order to elucidate and justify new approaches to social research (Aram and Salipante 2003; Baert 2003, 2005; Bosch 2007; Feilzer 2010; Fendt et al. 2007; Johnson and Onwuegbuzie 2004; Morgan 2007; Starkey, Hatchuel, and Tempest 2009; Watson 1997, 2011). It is argued that since disciplinary paradigms orient attention to certain “ways of seeing”, prioritise certain questions and research methods, and assert the dominance of certain theoretical traditions, they also constrain intellectual curiosity and creativity, and are blind to the multi-faceted nature of many if not most sociological phenomena (Feilzer 2010; Watson 1997). For instance, economics seeks a science of human action, at both macro and micro levels, underpinned by a paradigmatic adherence to “rational choice theory and the assumption of individual rationality as the arbiter of human judgment and incentive for action” (Starkey et al. 2009, p 550). But the application of the assumptions of economic rationality to complex human and social phenomena leaves too much out:

It exaggerates self-interest but ignores how self-interest interacts with the complex institutions of modern economic, social and political life. It is in this interaction through formal regulation and implicit codes of conduct, reputation, mutual co-ordination, instincts and structures of co-operation, that social life is created and sustained. (*excerpt from Starkey et al. 2009, p 551*)

In other words, the monocled perspectives of disciplinary methodologies are unable to perceive, let alone account for the complexity of an experiential world which is made up of many different elements or layers (Feilzer 2010; Watson 1997). To put it in terms of Latour’s metaphor, if everything is an effect of actor-networks then to study how things are produced, how problems are solved, and how worlds are made is to follow or trace how heterogeneous bits of the world become chained together (Latour 2005, 2007). In this view, “the social” is no longer understood as a particular domain of things among other non-social domains (e.g. as distinct from the “natural”, the “economic”, the “biological”, the “legal”, or the “psychological”), but should instead be understood as the making of “connections, associations, collections, whatever the name, between all sort [sic] of heterogeneous domains, none of them being “social” in the first meaning of the word... social is not the name

of any one link in a chain, nor even that of the chain, but it is that of the chaining itself” (Latour 2007, pp 3-4). Since there is no guarantee that this activity of chaining will be limited to forming connections between just legal things, or just economic things, or just political things, if we wish to gain a “fuller understanding” of phenomena, or to grasp the complexity of a real-world problem, it is unlikely that we will be able to do so from the confines of a single academic discipline or paradigm (Nowotny, Scott, and Gibbons 2003; Starkey et al. 2009; Watson 1997). The last decade has therefore seen increasing calls for researchers to side step the dichotomous choice of having to align oneself, a priori, with either one or other side of various longstanding philosophical dualisms<sup>3</sup> when framing research methodology, and instead adopt an approach of pragmatic pluralism (Aram and Salipante 2003; Feilzer 2010; Fendt et al. 2007; Johnson and Onwuegbuzie 2004; Morgan 2007; Starkey et al. 2009; Watson 1997).

### **So we should all be pragmatists?**

Described in abstract terms, as above, both pragmatism and ANT appear to quite deliberately bridge the key positions of the traditional paradigms (i.e. realism/positivism vs constructivism/interpretivism), adopting what is useful and leaving what is problematic from each. But what is the practical value of the pragmatist/ANT perspective as a methodology for the production of knowledge in existing disciplines? Can pragmatism or ANT actually be said to stand as alternatives to traditional disciplinary paradigms and methodologies? That is, is it really practical to expect the positivists and the interpretivists (or the adherents of any paradigmatic position for that matter) to suddenly admit the fundamental inconsistencies and flaws of their approaches (Baert 2003) and consequently become good pragmatists? I think the answer to this question is a clear “No!”, but not because the pragmatists are wrong. Rather, and somewhat ironically, it has more to do with the fact that, in most cases, such a switch of epistemological allegiance would make practically no difference to the production of knowledge in the various disciplinary projects.

Take, for instance, the pragmatic maxim that truth is never absolute, that knowledge is always provisional, that we must acknowledge that there is always the possibility that what we understand to be true may have to be revised in the future. Or take the ANT maxim that knowledge (science) only advances through controversy. Both of these statements are true enough as generalisations (else, indeed, they could not be taken as maxims). But what is the

---

<sup>3</sup> i.e. realism vs constructivism, free will vs determinism, positivism vs interpretivism, reductionism vs holism, etc.



practical relevance and significance of these maxims with respect to disciplinary inquiry? After all, even if one accepts that there are no absolutes, this does not require one to hold the established bodies of knowledge in natural and social science disciplines as uncertain or unreliable in any significant sense. As Constant noted, engineers and scientists “behave as though veridical, spatiotemporally universal knowledge does exist, and much more often than not they get away with it” (1999, p 327).

The point is that if all science is plagued by uncertainty and controversy, then it is hardly ever at the paradigmatic level. Rather, it is found primarily in the form of resistances encountered in the mangle of the practice of scientific research, which challenge and throw into question the particularities of method, e.g. data collection, measurement, analysis, etc. (Kuhn 1961; Pickering 1993, 1994). Such operational difficulties may be significant for the individual researcher, but they are unlikely to amount to a fundamental challenge to the disciplinary paradigm in which the researcher is embedded. This does not, of course, rule out the possibility that controversies might persist and ultimately be found to be irresolvable within a given paradigm. Indeed, as Kuhn (1962) famously demonstrated, science does not stick rigidly to any one paradigm, but rather moves through periods of paradigm stability and controversy, shifting from one dominant paradigm to another as new theories and evidence emerge.

Thus, deep within disciplinary domains, in context of the interests, aims, and objectives of traditional disciplinary inquiry, pragmatism would appear to have little to add. Scientists of all persuasions already always approach their various objects of interest with certain preconceptions, and, through various processes and at various scales (both transtemporal and transpersonal), those preconceptions are already always open to question. This would seem to suggest that in the context of normal science (Funtowicz and Ravetz 1993) or basic research (Hadorn et al. 2008) the “practical consequences” of becoming a pragmatist would be trivial.

Rather, where pragmatism takes on true significance as an alternative to extant disciplinary approaches is in contexts of cross- or transdisciplinary inquiry. Two situations come to mind. The first is where the objectives of inquiry call for deliberately crossing disciplinary boundaries, perhaps to generate a “fuller understanding” of phenomena (Starkey et al. 2009; Watson 1997, 2011), or to seek out solutions to complex real-world problems (Fendt et al. 2007; Hadorn et al. 2008; Klein et al. 2001; Tranfield and Starkey 1998; van Aken 2005; Zierhofer and Burger 2007). The second situation, which may go hand in hand with, or indeed lead to the first, is where there is significant uncertainty about how to fulfil the cognitive interests of inquiry (Alferoff and Knights 2009; Funtowicz and Ravetz 1993;

Nowotny, Scott, and Gibbons 2001; Baert 2005). That is, where the inquirer is confronted with strategic<sup>4</sup> uncertainty about the nature of the research problem and the object(s) of interest, about what features of the empirical domain might be relevant and significant, and about what theory and methods might be appropriate.

In these situations the pragmatic approach is attractive because it eschews a priori constraint on which theories and methods might prove useful in understanding the problem or phenomena under investigation (Baert 2005; Bosch 2007; Feilzer 2010; Watson 1997). Rather, the approach is pluralist, allowing for the selection and mobilisation of empirical and theoretical material according to what the researcher judges to be relevant, plausible, and significant with respect to the subject under consideration (Bosch 2007). It is effectively an approach whereby the researcher “draws elements from various disciplines or perspectives to produce what amounts to their personal paradigm – with its own ontological, epistemological and methodological integrity – to stand as the conceptual foundation of that particular piece of research” (Watson 1997, p 6).

## **Performing the role of Chief Risk Officer: a pragmatic inquiry**

Since about the early 1990s research in the fields of organisation studies and management has been motivated toward transdisciplinary approaches in response to a perceived lack of relevance of traditional academic research for management practice (Fendt et al. 2007; Huff and Huff 2001; Shani et al. 2008; Starkey and Madan 2001; Starkey et al. 2009; Tranfield and Starkey 1998). Over the latter half of the twentieth century the strategic direction of management research became increasingly determined by the interests of the academic community with little consideration given to the practical interests of managers, such that business schools were increasingly criticised for producing knowledge of little or no relevance to managers (Ghoshal 2005; HEFCE 1998; Hoffman 2004; Porter and McKibbin 1988). The subsequent “relevance” debate opened up the question of what the strategic mission of business schools should be (Huff and Huff 2001; Starkey and Madan 2001; Tranfield and Starkey 1998). Although the pure academic project was not deniable, the idea that the production of knowledge in management should imitate the reductionist model of the physical sciences was challenged (Fendt et al. 2007; Starkey et al. 2009; Starkey and Madan

---

<sup>4</sup> Refer to the discussion in Chapter 6.

2001; Tranfield and Starkey 1998; van Aken 2005). Critical questions were raised about the continued legitimacy of business schools if they failed to re-engage with managers and support management practice (Starkey and Madan 2001).

The perceived need for relevance motivated a broad, multi-faceted reconceptualisation of management research in the 1990s and 2000s. Various related thematic strands are identifiable, including: promoting a transition to “Mode 2” research (Huff and Huff 2001; MacLean, MacIntosh, and Grant 2002; Starkey and Madan 2001); building collaborative, interventionist approaches between academics and practitioners (Adler, Shani, and Styhre 2004; Hatchuel 2000, 2001; Shani et al. 2008); and positioning management as a “design science” equivalent to engineering in the physical sciences or medicine in the biological sciences (Fendt et al. 2007; Pandza and Thorpe 2010; Starkey et al. 2009; Tranfield and Starkey 1998; van Aken 2005). In this regard, the label of “transdisciplinary”, in management research as elsewhere, does not refer to a single methodological approach, but rather encompasses a plurality of approaches and research forms (Nowotny et al. 2003; Zierhofer and Burger 2007).

A common feature of the discourse around the future of management research is the idea that achieving practical relevance requires researchers to re-engage with the world of management practice. Researchers are encouraged not just to enter into the real-world contexts in which managerial activity takes place (which they have always done to varying extents), but to orient their attention to the problems of management practitioners and to open up to the complexity of those performances (Aram and Salipante 2003; Fendt et al. 2007; Starkey et al. 2009; Watson 1997, 2011). However, such engagement renders methodological choices problematic.

First, the involvement of practitioners and other non-academic actors in defining the research problem throws into question whose interests are to be served and what is to be achieved by the project (Funtowicz and Ravetz 1993; Hadorn et al. 2008; Nowotny et al. 2003). Second, to the extent that academics and non-academics become entangled as co-producers of knowledge this throws into question the notion of academic distance or objectivity. Research methods can no longer be seen as innocently capturing and representing reality, but must be understood as involved in “ontological politics” – the making of social realities (Latour 1998; Law and Urry 2004). Third, speaking from experience, opening up to the complexity of real-world performances, and indeed becoming involved in them, will likely make it difficult to judge the ontological nature and scope of the research object before the fact. Law and Urry (2004) call attention to a fourth difficulty, that current research methods do not resonate well with the complexity of “reality enactments”:

They deal, for instance, poorly with the *fleeting* – that which is here today and gone tomorrow, only to reappear the day after tomorrow. They deal poorly with the *distributed* – that is to be found here and there but not in between – or that which slips and slides between one place and another. They deal poorly with the *multiple* – that which takes different shapes in different places. They deal poorly with the non-causal, the chaotic, the complex. (*excerpt from Law and Urry 2004, p 403*)

Thus, in transdisciplinary settings the methodological relationship between cognitive interests, theory, methods, and reality may not be well understood (Baert 2005). There is then a need for social science to re-imagine itself, to find ways of coping with and engaging in a world where “social relations appear increasingly complex, elusive, ephemeral, and unpredictable... perhaps, for instance, there is need for ‘messy’ methods” (Law and Urry 2004, p 390).

### **Shadowing as a way to fulfil a practical interest (and mitigate risk)**

I initially gained access to the Watercare organisation in 2004 as an embedded researcher on an ultimately unsuccessful project to investigate the concept of “sustainability” with respect to the planning of urban water and wastewater systems (this background is explained in more detail in Appendix II). As a result of this background I was reasonably familiar with the Watercare organisation and its planning processes, and I had even acquired some practical experience with designing frameworks for risk assessment and the evaluation of projects and programmes for capital decision making. In early 2007 I recognised an opportunity for research centred around the new Corporate Risk Manager’s project to redevelop Watercare’s Risk Management functions. That opportunity was framed out of a translation of the objectives and understandings acquired during the earlier project, central to which was a practical interest. That is, the funding for the original project had been conditional on the research producing a practical output, either a technological development, or improvements to business planning processes around the firm’s core technology programme. Broadly, I reasoned that the CRM’s project would open up a core component of Watercare’s capital decision making infrastructure, thus providing an opportunity to study how that infrastructure was constituted. Since risk was important for justifying expenditure and framing problem-solving behaviour within Watercare, it could be expected that as risk concepts and practices evolved, this would eventually influence the objectives, form and content of capital projects, and, hence, the long-term performance of the water and wastewater infrastructure. If I could document and analyse that evolution then the resulting insights might productively inform governance and strategic planning processes both within

Watercare and the broader socio-technical system of which the firm was a part.

This initial framing of the project was, however, fairly rudimentary. I had identified where and roughly what to look at through a crude functional analysis of the Watercare system, but I did not know before hand what the project was or even whether I would be successful in framing an academically significant inquiry in that context. In this regard, my research proceeded from an initially precarious position. The design of Watercare's capital planning processes and decision infrastructure constituted the broad "area of concern" but I remained uncertain about the specific problem to be investigated, about the scope of the "framework of ideas" relevant to that area of concern, and hence also about the appropriate methodological stance (Jackson 2000). The decision to shadow Watercare's Corporate Risk Manager was essentially a strategy aimed at mitigating the risk of that starting point. That is, I clearly needed to be in the places where the "opening up" of Watercare's Risk Management functions would occur, and since the CRM intended to initiate and co-ordinate important changes in the system, it was intuitive that his role was an important vector to be followed. But since I was not sure what specifically I was looking for I intuitively set out to collect as much "data" as possible, and figure out how to analyse it later. An ethnographic style of data collection was therefore intuitive.

Shadowing "is a research technique which involves a researcher closely following a member of an organisation over an extended period of time" (McDonald 2005, p 456). It has been employed as a method for ethnographical studies of the behaviours of individuals and groups in a range of disciplines; including cognitive science (e.g. Hutchins 1995), education (e.g. Polite, McClure, and Rollie 1997; Wolcott 2003), engineering design (e.g. McGarry 2005), information studies (e.g. Hirsh 1999), nursing (e.g. Vukic and Keddy 2002), organisation and management (e.g. Bonazzi 1998; Mintzberg 1970; Noël 1989; Snyder and Glueck 1980), and sociology (e.g. Fenton et al. 1997). Shadowing has been performed in different ways in different contexts (i.e. studies of varying length, location, intensity, and structure), and is often combined with other techniques, including in-depth interviews, diary studies, surveys, documentary analysis (Czarniawska 2008; McDonald 2005). Supplemental techniques are employed to overcome what Czarniawska (2008) called the problem of simultaneity – that the sites where managing or organising takes place may not always be where the researcher is, seemingly requiring the researcher to be in more than one place at once. Supplemental techniques allow the researcher access to what happened in those other sites through, for instance second hand accounts or via formal and informal representations of actions (McDonald 2005).

In her review, McDonald (2005) identified three traditions of shadowing research. The

first is as a method of vocational education whereby the student can gain first hand experience of professional practice (e.g. Paskiewicz 2002). The second and most common form is where the method is employed to record a detailed log of the minute-to-minute actions of the individual(s) being observed, utilising quantitative or mixed quantitative/qualitative frameworks (e.g. Perlow 1998; Walker, Guest, and Turner 1956). In both cases the focus is on the individual, rather than the department, company or function, and the objective is reveal the direct, first-hand nature of the experience of performing the role in which that individual is engaged (McDonald 2005). The third form of shadowing also embodies concern for behaviour and experience, but this is secondary to the objective of revealing the point of view and rationale from which the role is approached (McDonald 2005). The difference is that the research is primarily interested in the programme of an individual's actions, rather than the peculiarities of the immediate performance, and thus seeks to surface patterns of purpose and meaning from a grounded, qualitative study of activities. McDonald (2005) noted that this is the rarest form of shadowing, but also that with the most potential for extending the reach of current organisational research; a view echoed by Watson (2011) who argued that organisational ethnography should focus on investigating "how things work" in organisations rather than being distracted by concerns with capturing the subjective, lived experiences of organisational actors.

Although I adopted the shadowing method intuitively rather than as the result of a formal research design, my performance of it was consistent with the latter form identified by McDonald (2005). I shadowed Watercare's Corporate Risk Manager for approximately 15 months (March 2007 to May 2008). The main sources of data were observations and transcriptions of meetings and workshops that the CRM attended, and monthly semi-structured one-on-one interview/discussion sessions between myself and the CRM where I asked him to reflect on his approach, problems and issues he was encountering, what was successful, and what he was learning (the Dialogues are appended on pages 237-336). Since my workstation was located adjacent to the CRM's desk in Watercare's corporate office, there were also frequent impromptu discussions on various issues that arose on a day-to-day basis. The CRM also gave me open access to documents that he prepared (e.g. emails, reports, memos, discussion documents etc.), as well as books and articles that he found influential. From the outset I was rather more interested in how the CRM's programme of actions would produce changes in the company's decision making infrastructure than with minutiae of his performance, or his personal or emotive experience of the role.

It also became apparent from a review of the literature that the role of the Chief Risk Officer and the implementation of Enterprise Risk Management were relatively virgin areas of

academic endeavour (this gap was framed up in Chapter 1). Since I was shadowing a CRO engaged in the implementation of ERM, I realised that I was in a uniquely privileged position to contribute to this gap, and I therefore shifted my attention to focus specifically on the CRM's role and motivating rationale. While I remained uncertain for some time about what specific problem the thesis would eventually address, I continued to intuitively focus on the functional questions of what the CRM was doing, how, and why; with those questions being directed at the overall programme of his actions and the broader rationale on which those actions were based, rather than the contingencies of the moment-to-moment performance. The sources of data identified above reflect this broader focus.

### **Two projects: collaboratively seeking understanding of practice**

The CRM was not just new to his job, but also to the role, and with little prior experience or formal training as a CRO he was effectively “learning the ropes” of corporate risk management. At least initially the CRM's project was the (re)design of risk frameworks, registers, and models, and he expressed a relative confidence in his approach to this project. It was only later, after reflecting on how the various problems and dilemmas that he encountered were challenging that approach, that the CRM became more broadly concerned with the question of what the best approach to operationalising ERM might be, and the question of how best to fulfil the decision support function of his role in particular. The various questions that framed the CRM's project in this regard are identified in the left-hand column of Table 2.1 (progression top to bottom reflects the evolution of the CRM's project over time).

Since I began the research knowing relatively little about ERM and the role of the CRO, I had to acquire or otherwise develop a certain knowledge base in order to examine and think through the empirical moments and to formulate a conceptual framework to structure those insights. Given that the CRO role and the “how to” of ERM implementation are relatively virgin areas academic endeavour it is possible that I could have eschewed a theoretical search in favour of a completely grounded study, since theoretical naivety is a deliberate methodological stance for such studies (Creswell 1998; Glaser and Strauss 1967 [2006]). However, at the time I was shadowing the CRM I was uncomfortable with such an approach because I did not have a firm grasp of what questions the research would eventually address. I felt it was therefore necessary to engage with the relevant literatures in order to frame up those questions in parallel with shadowing the CRM. My project therefore evolved in two intertwined parts: first to shadow the CRM and make sense of what was going on, and second

to work out the relevant theoretical framework and particular problem that the research was to address. In addition to collecting empirical data I explored various literatures around the topics of risk, risk management, enterprise risk management, decision making, and knowledge management; guided primarily by what I felt was relevant for understanding the CRM's performance (see Box 2.1). Through reflection upon the various problems and dilemmas that I encountered, I also became concerned with the question of what the best approach to operationalising ERM might be, and the question of how CRO's can best fulfil the decision support function of their role in particular. The various questions that framed my project in this regard are identified in the right-hand column of Table 2.1 (progression top to bottom reflects the evolution of my project over time).

*Table 2.1. Parallel streams of inquiry*

<b>Practitioner stream</b> <i>Questions framing the CRM's practice</i>	<b>Academic stream</b> <i>Questions framing my research and the thesis</i>
<p>How can the corporate objectives be deconstructed into a comprehensive and objective risk hierarchy?</p> <p>How can that hierarchy be organised so as to be able to objectively quantify and aggregate risks at various levels?</p> <p>What do actors worry about (i.e. what matters?) at different levels within the organisation?</p> <p>How can the corporate objectives be married with operational criteria within the risk framework to reflect what matters?</p> <p>What does "strategic risk" mean in the Watercare context?</p> <p>What is the purpose of "risk data" in a mature organisation? How does "risk data" support decision making?</p> <p>What is the value of making explicit what we already know in a mature organisation? (i.e. what is the value of formal risk assessment?)</p> <p>What is the appropriate level of detail for risk frameworks, risk assessments, and risk data?</p> <p>What does "integrated Risk Management" mean in practice? How can abstract ERM models be integrated into existing business processes?</p>	<p><u>Empirical Questions</u></p> <p>What did the CRM want to do at Watercare, and why? Were his proposals reasonable? Were they justified?</p> <p>What happened when the CRM intervened in the organisation, and why? What can be learned about designing ERM infrastructure from those experiences?</p> <p>What were the key problems and dilemmas that the CRM encountered, and how did he resolve them, if at all? How do the CRM's experience problematise ERM more generally?</p> <p><u>Theoretical Questions</u></p> <p>What is Risk? What are the implications of different disciplinary perspectives on Risk? What is the relationship between Risk and Uncertainty?</p> <p>What is Risk Management? What are the implications of "Risk Management = Uncertainty Management" with respect to identifying and describing Risk Management as a distinct category of processes and practices within the organisation?</p> <p>If Risk is, in practice, knowledge about the future, then what is the relationship between Risk Management, decision making, and knowledge management? What is the value of making Risk explicit?</p> <p>How should CROs approach the task of organising ERM in organisations?</p>



**Box 2.1.** *Relevance of different literatures to understanding Risk and Risk Management*

- Literature on ERM frameworks (e.g. COSO 2004; AS/NZS:4360 2004a) and capability development (e.g. Lam 2003; MacGillivray, Hamilton, Strutt et al. 2006; MacGillivray et al. 2007a; Shenkir and Walker 2006; Short and Clarke 1992; Ward 2005), and on the role of the Chief Risk Officer (e.g. Aabo et al. 2005; Power 2005b; Ward 2001), was relevant because the CRM was attempting to implement ERM in the organisation.
- Literature on engineering risk and modelling systems reliability (e.g. Ayyub 2003; Keey 2000; Moubray 1997) was relevant because Watercare employed such models to inform the scheduling of routine maintenance activities.
- Literature on the meta-theory of risk (e.g. Althaus 2005; Fischhoff, Watson, and Hope 1984; Kaplan and Garrick 1981; Macgill and Siu 2004; Renn 1992, 1998; Rescher 1983; Rosa 1998) was relevant because I wanted to understand what risk is.
- Literature on the social construction of risk (e.g. Douglas 1985; Hilgartner 1992; Kasperson and Kasperson 1996; Krimsky and Golding 1992) was relevant because there was clearly a significant degree of subjectivity involved in assessing risk in various contexts within the company.
- Literature on risk, decision making, and rationality (e.g. Aven 2004; Cabantous and Gond 2006; Hansson 1994; Jaeger et al. 2001; Kleindorfer, Kunreuther, and Schoemaker 1993; Laroche 1995; March 1988a) was relevant because risk was an important variable in capital decision making, there were clear similarities between the normative Risk Management procedure and the traditional decision making process, and decision support was an explicit function of ERM.
- Literature on the production and management of knowledge (e.g. Alvesson and Kärreman 2001; Buenaño 1999; Tsoukas and Vladimirov 2001) was relevant because risk could be seen as just a label for what actors knew about the future in more or less detail.
- Literature on classification (e.g. Bowker and Star 1999) and the role of boundary infrastructures in moderating communication between communities of practice (e.g. Brown and Duguid 1991, 2002) within the organisation was relevant because this appeared to be a primary role of the risk framework.
- Literature on the performativity of objects (e.g. Law and Singleton 2000, 2003) was relevant because risk could also be seen as something not “out there” waiting to be discovered, but as something performed into being by particular arrangements of people and infrastructure (e.g. Hilgartner 1992; Scott and Perry 2006). Sociological perspectives on the organisation of risk management (e.g. Hutter and Power 2005b; Miller, Kurunmäki, and O'Leary 2006; Power 2005b, 2007) and on decision making (e.g. Callon 1998c, 1999; Chia 1994; Cooper 1986) were relevant for the same reason.

Czarniawska (2008) has noted that shadowing is an interactive if not collaborative mode of engagement with the world:

Shadowing creates a peculiar twosome – the person shadowed and the person doing the shadowing – in which the dynamics of cognition become complex and therefore interesting. There is a mutual observation, an establishing of similarities and differences; then there is a focus created by the movements of the person shadowed, and the double perception of a kind – the researcher guesses (and asks about) perceptions of the events being perceived as well. (pp 10-11)

Bosch (2007) has noted that social science is also characterised by a reverse process not present in the natural sciences, which he called a “double hermeneutic”, referring to “the interpretation by the social scientist of social phenomena, and the subsequent interpretation of findings of the social sciences by social actors” (p 193). This is easily understood as a temporally separate movement, where social phenomena are studied by social scientists who then publish their results, and those published results are then reinterpreted by other social actors in another time and place removed from the original study. In this way social science can be seen to be constitutive of the very phenomena it seeks to study (Callon 1998b; Law and Urry 2004; Osbourne and Rose 1999). A methodologically significant feature of my shadowing of Watercare’s CRM was our co-performance of this “double perception and hermeneutic” in real time due to the fact that the CRM and I were participants in each other’s projects. That is, while the CRM and I were clearly working on our own respective practical and academic projects, those projects shared a common interest and were collaboratively interactive. I was as much engaged in perceiving, interpreting, and trying to make sense of the CRM’s performance with respect to his overarching rationale (this being the primary focus of the CRM’s project), as the CRM was engaged in thinking through or theorising how to approach the CRO role (this being the primary focus of my project).

In some cases I gained first hand experience of the practical side of ERM implementation by performing work for and with the CRM, such as with the Strategic Risk Assessment (described in Chapter 4) where I collated managers’ responses to the CRM’s survey, worked with the CRM to mould those responses into the Strategic Risk Framework, and helped prepare the workshop presentation. More frequently, however, we collaborated through discussion where we considered various issues and dilemmas, explored ill-defined problems and conceptual conflicts, and hashed out ideas about how to proceed. In these discussions I was able to bring a theory-based perspective to bear on the problems that the CRM encountered in practice. The CRM benefitted from having a sounding board – indeed, he commented that I was often the only person in the company who was both interested and

able to understand the problems he was grappling with – and from access that I provided to alternative perspectives, new insights, and theoretical models, drawn from a range of disciplinary sources, which he might not otherwise have had the time or resources to acquire. I benefitted from the resulting insights into the CRM's thinking about his role. A critical characteristic of the CRM's practice in this regard was the degree to which he reflected on his approach, about the efficacy of his interventions, and about the various problems and dilemmas that he encountered. This rigour was, I believe, a product of the CRM's disposition (he openly admitted that he often spent too long thinking things through rather than getting on with doing them), and his own academic training (PhD). I was able to gain access to the CRM's reflections through our discussions and these became an indispensable component of the empirical data set. The regular monthly sessions were particularly important as opportunities (usually one to two hours in length) for the CRM to step back from and reflect on his work.

Thus, during the shadowing phase of the inquiry the CRM and I were not so much collaborating on a common project as working on two complementary projects, one practice-oriented with reflection on theory (the CRM's project), the other theory-focused with reflection on practical implications (my project). Those two projects were convergent around a common practical intent (i.e. how to intervene in the organisation to support decision making), and were mutually and productively collaborative.

### **Pragmatic hermeneutics in response to strategic uncertainty**

Shortly after Dialogue 14 was recorded in May 2008 I departed Watercare to concentrate on analysing and documenting the data I had already collected over the preceding 15 months. That decision was not based on knowing that I had collected sufficient information to address my research questions (i.e. saturation). Indeed, even by that point I had only a vague notion of what those questions were, and what arguments the thesis might mount in respect of them. But fifteen months of shadowing the CRM had produced a substantial volume of data which would take time to properly analyse and understand. Based on a preliminary parsing of that information, and my exposure to various literatures, I had loosely framed a number of key ideas around the decision support function of ERM as a central theme. In that regard, the CRM's paradigm shift (Chapter 5) suggested a convenient break point defining a potential core around which to structure the thesis (i.e. interpretation, implementation, re-interpretation). But I realised that those ideas were not narrow and would take time to develop into a written thesis, and that continued engagement with the practice context

would distract me from that task. While I did keep in occasional contact with the CRM after that, I did not formally collect any further information about his experiences at Watercare.

The status of my inquiry at the time of my departure from Watercare points up the other constitutive driver of the research methodology – that is, the degree of uncertainty which characterised my project, as well as that of the CRM. The sense in which I am talking about uncertainty here is strategic. That is, uncertainty about the nature of the problem (or indeed whether there is a problem), about what questions to ask, and about what is relevant and significant, both empirically and theoretically. This was not just a state of mind which I experienced in the early stages of the inquiry when I was grasping for a project. Rather, it was something which persisted for both the CRM and I even after we had conceptualised what it was we were trying to do in our respective projects.

This was because the very natures of the objects of our attention (i.e. risk, risk management, ERM, and even the concepts of decision and decision support) were called into question through our respective and shared encounters with theory and practice. This can be seen most clearly in our shift in thinking about the nature of Risk. Initially, in Chapter 3, Risk is something objective, a quantifiable property of the real world that can and should be measured and calculated through the correct scientific procedures. In Chapter 4, Risk becomes something both real and subjective, a product of social value judgements as well as of the natural variability of phenomena. Then in Chapter 5 Risk is no longer a specific state of knowledge, but merely a category of more or less specific perceptions about the future which serve as points of attention for managerial activity. The problems, challenges, and dilemmas which gave rise to these shifts in conception were not reconcilable as mere issues of discrepancy between our theoretical preconceptions of what Risk should be and the mangle of organisational practice (Pickering 1993, 1994). The subjectivity that we encountered was not merely an “error”, arising from practical resistances to the calculation and communication of Risk, which had to be overcome. Nor was it the product of our methods of engagement with those practices (i.e. merely perceived rather than actual). It resisted all attempts to make it go away, or to make it irrelevant, such that we eventually had to acknowledge that “subjective” was a real, defining characteristic, in both ontological and epistemological senses, of this thing called “Risk” in the Watercare organisation. This was far beyond the kind of “re-thinking” which characterises all research to varying extents, e.g. revising one’s calculations to accommodate new data, or modifying methods to accommodate unforeseen resistances (Pickering 1993, 1994). Rather, both of us were ultimately forced to admit that the world (or in this case, Risk) is (or might be) a fundamentally different kind of place (object) than we had originally assumed to be.

In this regard, both of our projects were characteristic of the pragmatic approach in terms of questioning and re-articulating preconceptions through dialogical confrontation (Aram and Salipante 2003; Baert 2003, 2005; Bosch 2007). The principle behind the dialogic method stems from the recognition that while theoretical preconception is a necessary precondition for knowledge production, our subjective interpretive schemas also limit what we can know about the world, but can never be transcended in the sense of obtaining a completely objective or neutral vantage point (Baert 2003, 2005; Bosch 2007; Rescher 2005). In contexts of transdisciplinary inquiry, where we are faced with the uncertainty which attends the crossing of disciplinary boundaries, the only path that remains open to us is to improve and evolve our preconceptions by accommodating alternative perspectives (Baert 2003, 2005; Bosch 2007):

This sensitivity can be achieved through a conscious effort to remain open to other traditions and learn from them. This openness and willingness to learn from other traditions is central to the way in which the dialogical model can be employed in the philosophy of the social sciences. (*excerpt from Baert 2005, p 196*)

The pragmatic approach therefore conceives of inquiry as an active process of hermeneutic dialogue: we must start from somewhere, initially gaining access to what is being studied through our preconceptions, but then proceeding toward a fuller and more truthful (warranted) understanding by reassessing those views in the light of novel experiences and encounters with difference (Aram and Salipante 2003; Baert 2003, 2005; Bosch 2007; Rescher 2005). Figure 2.1 (page 59) borrows the words of others to conceptualise the process as an iterative cycle – the hermeneutic circle – between contextualisation and abstraction, seeking a synthesis between the parts and the whole by constantly generating, contrasting, and comparing empirical and theoretical concepts (Aram and Salipante 2003; Bosch 2007; see also Klein and Myers 1999).

It is in these terms of pragmatic hermeneutic dialogue that the performance of my inquiry should be understood. Figure 2.2 (page 60) conceptualises the inquiry as a number of parallel but interactive and mutually reflexive work streams, constituting a hermeneutic cycle of interpreting, making sense of, and theorising the practice of ERM implementation (Figure 2.2 shows the relationships between the work streams and the questions which guided them):

- 1) Collecting data: the work of shadowing the CRM, described earlier. (A notable aspect of the data collection process was the formation of first and general impressions through the experiences of being involved in particular empirical moments. These

impressions were often, but not always, a good guide as to what was significant about a particular situation, whether in its own right or in relation to the broader inquiry.)

- 2) Exploring literature: reading literature across a range of disciplines, guided by a sense of what was relevant with respect to the CRM's performance, and the concept of decision support.
- 3) Sorting out the empirical data: describing, summarising, and providing contextual explanations of what happened in various empirical moments.
- 4) Examining the empirical data: working out how each empirical moment was relevant and significant with respect to the CRM's project and rationale, whether individually or collectively. Relevance and significance was judged with respect to process (i.e. how the CRM did what he did) or outcome (i.e. whether the CRM was successful or not).
- 5) Theorising the concept of Decision Support: seeking an understanding of decisions and how people make them in practice. This aspect of the inquiry evolved from certain starting points, which came out of (2) above, to focus mainly on how the relationship between knowledge and action varies with context as a basis for formulating decision support strategies in those contexts.
- 6) Theorising empirical insights: bringing in theory to explain and characterise empirical moments (e.g. using the concept of knowledge facilitation to explain and characterise the CRM's strategic risk assessment work), and relating key themes or threads from the empirical narrative to corresponding themes or constructs in the literature (e.g. relating the basis of the CRM's paradigm shift to changing assumptions about the nature of Risk).
- 7) Contextualising the Decision Support concept: drawing on insights from the analysis of the CRM's performance to theorise strategies for decision support in context of the CRO role in general (e.g. the dilemma posed by the fact that CROs are generalists not experts).

Some of the work streams were primarily empirical (i.e. concerned with collecting data and making sense of what was happening empirically), some were primarily theoretical (i.e. concerned with exploring the literature and understanding new concepts), and some were integrative (i.e. relating theory to practice and practice to theory). The analysis of the empirical data and of the literature began with the reading of texts; both empirical texts (i.e. transcripts, memos, reports, working papers, etc.), from which the central concepts and

themes of the CRM's performance emerged, and texts from the literature (i.e. books, journal articles, etc.), through which key theoretical concepts were identified. The integrative process can be thought of as one of breaking down the world into conceptual parts while simultaneously trying to grasp how those parts related together to constitute a complex whole<sup>5</sup>; thus reflecting an abductive logic (Fischer 2001). The writing of the Dialogues was an important part of the method, because it encouraged the identification and elaboration of key constructs (both empirical and theoretical), and reflection about the relationships between them. The movement between the work streams was iterative. That is, insights, understandings, and discoveries from one stream would influence work in the others, and this was often recursive as new developments would necessitate the rethinking of old. The following examples illustrate the process:

- I initially took direction from the CRM's expression of his motivating rationale; i.e. to support good decision making with good "risk data". Since I was following a Chief Risk Officer engaged in the implementation of Enterprise Risk Management it was logical to ask certain questions: What is ERM? What do CROs do? How do they do it? What do they need to know? My recognition of a corresponding normative agenda and "vision" of good Risk Management in the ERM literature, and the relative lack of theoretical guidance for its implementation, subsequently reinforced my attention to the CRM's rationale and his vision for developing ERM at Watercare. My exploration of the literature was, in turn, further oriented by the CRM's reflections about his approach, particularly in the latter half of 2007 and early 2008 when various resistances to his endeavours forced him to question his decision support rationale.
- Since the CRM was expressly concerned with Risk and Risk Management, it was logical to explore literature dealing with these objects. That review uncovered a diversity of disciplinary perspectives, all of which I perceived to be relevant, but individually none of which provided sufficient understanding to make sense of everything that seemed to be relevant in the domain in which I was located (see Box 2.1, page 51). Each of the literatures offered a different way of seeing the world and understanding what was going on when the CRM intervened in the Watercare organisation. I realised that these different perspectives would need to be accommodated and integrated in order to evaluate the implications of the CRM's

---

<sup>5</sup> Here the term "whole" can be understood to refer both to the empirical situation (i.e. the CRM's performance) and to my inquiry.

experiences with respect to his decision support rationale. In this regard I took direction from what I was observing: the CRM's encounters with the subjective side of Risk in his engagements with different communities of practice in the organisation; the relative absence of capabilities to calculate Risk in accordance with its theoretical specification; the looseness of the term "risk" in practical usage. The various inconsistencies and contradictions between these observations and the theoretical concepts from the literature prompted reflections and reinterpretations of the CRM's experiences from new perspectives.

- The CRM's decision support rationale and the apparent relationship between Risk Management and corporate decision making prompted inquiry into decision theory as well as knowledge and knowledge management. Over time, my evolving understanding of decision making provided a conceptual point of reference for interpreting empirical moments. For instance, in Dialogue 11, the CRM's puzzlement about the differences in the use of "risk data" in capital versus operational decision making in the company resonated with ideas (to which I gave voice at the time) about different forms of decision making and the appropriateness of formal Risk Assessment as a means of supporting decision making in those contexts. Concepts from the knowledge management literature reinforced the discomfort experienced in trying to tell engineers, who are experts in their fields, how they should be specifying and calculating Risk (i.e. via new risk frameworks). This discomfort persisted as a controversy that needed to be resolved, i.e. if CROs are not to tell others about Risk then what are they to do instead?



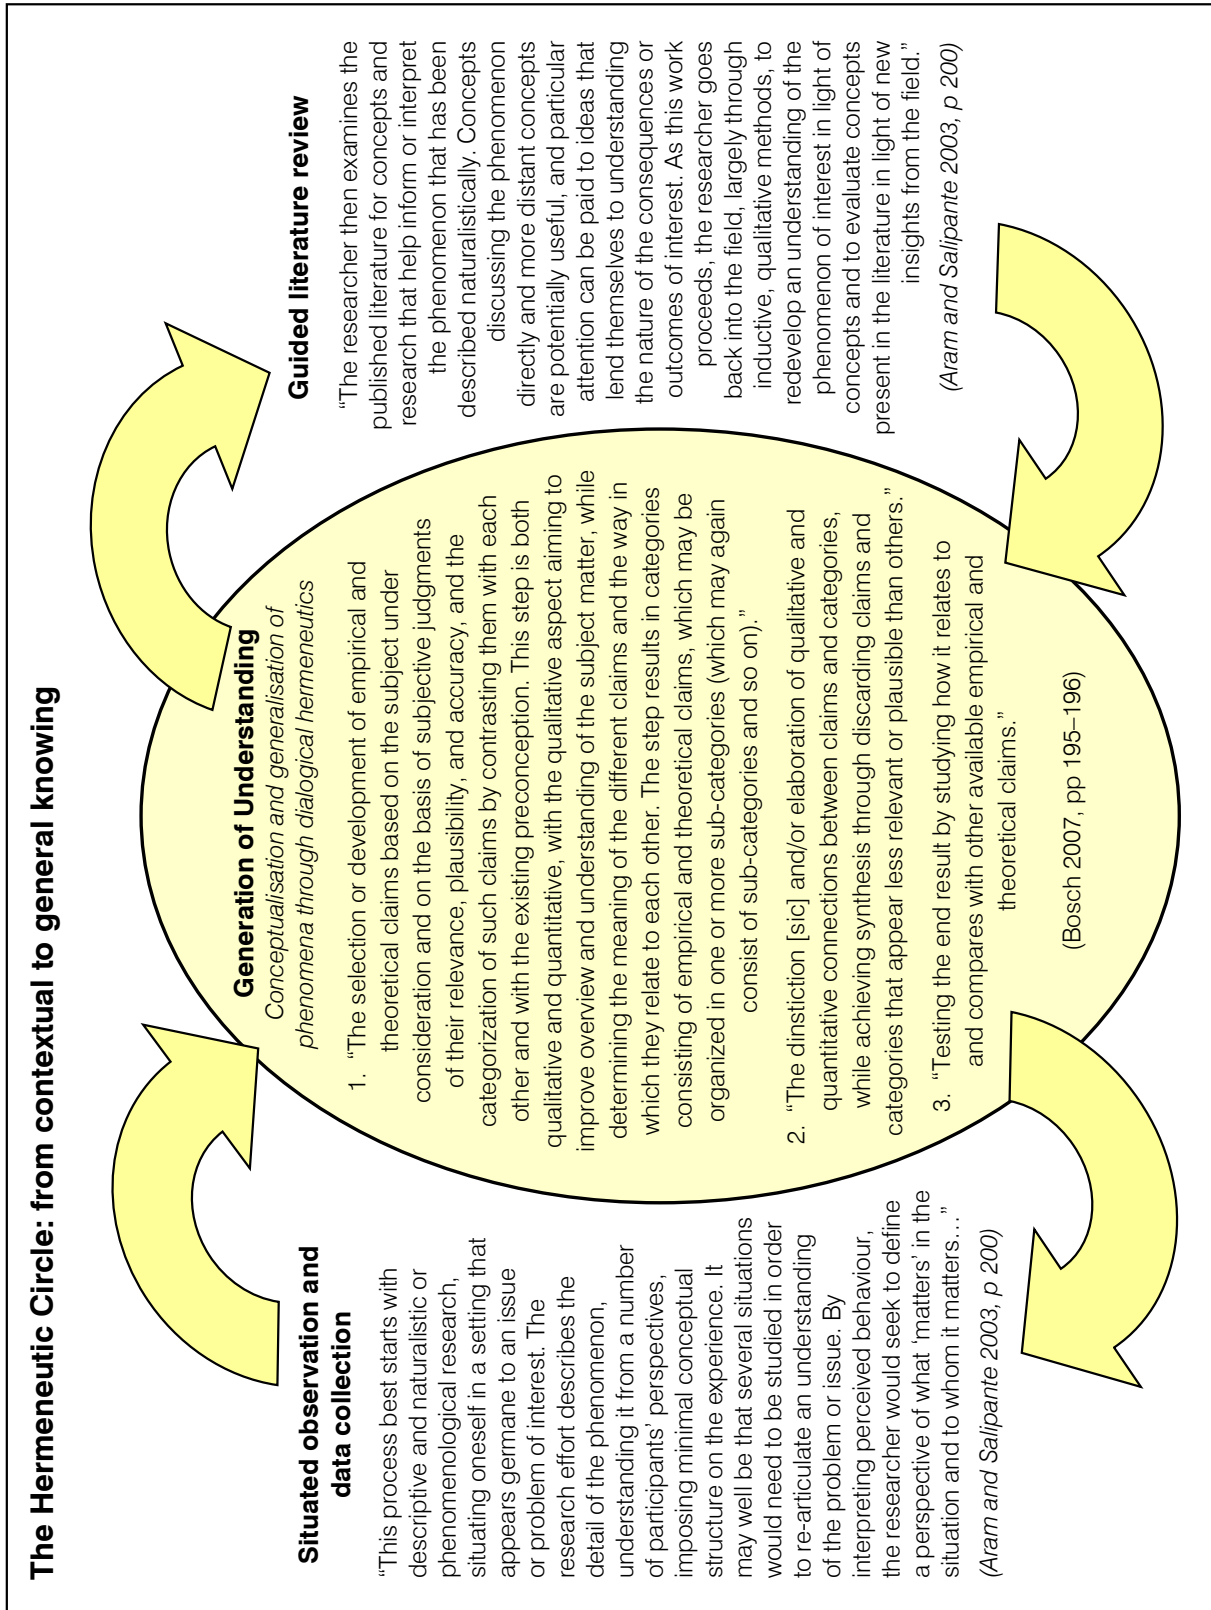


Figure 2.1. The Hermeneutic Circle

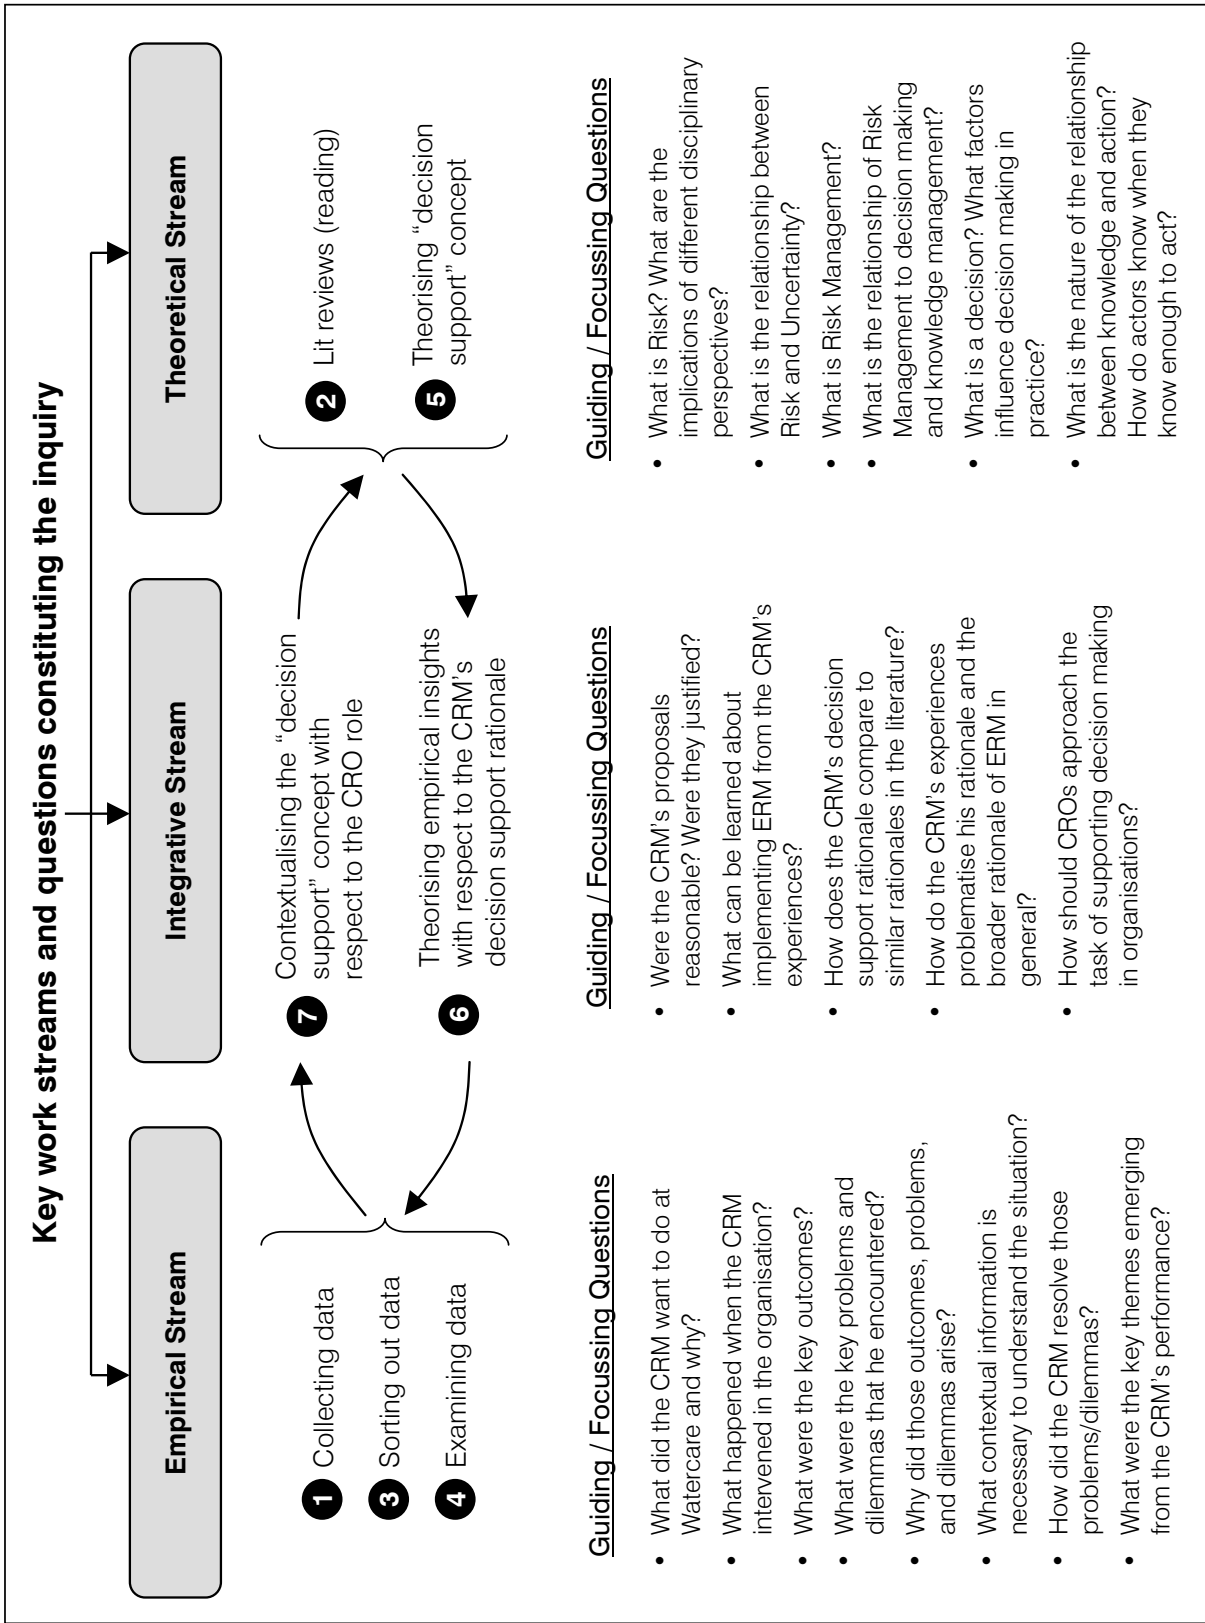


Figure 2.2. Key work streams and questions constituting the inquiry

## Judging the account: workability

Following the pragmatic principle that the truth value of a good theory lies in its problem-solving power, Watson (1997) suggests that the practical test of management theory should be that if someone:

...considered the theoretical interpretation of the nature of managerial work presented in the account of the research and then entered a managerial context and acted in a way informed by those interpretations, they would cope more successfully in that setting than if they followed interpretations based on poorer research – on accounts which, in this sense, are less true. This means that any social actor [entering that context] would be better placed when trying to ‘learn the ropes’ than someone who was informed by a less true research account. They would have greater power to act in relation to that environment than someone less truthfully informed. (excerpt from Watson 1997, pp 6-7)

I believe the contributions of this thesis are relevant to practising Chief Risk Officers – that, having read my account, CROs will be able to cope more successfully in their roles (the response of Watercare’s CRM to this thesis gives me confidence in this regard, as do the responses of others of the managerial persuasion who have read it). However, I cannot claim that my inquiry produced knowledge that was both practically relevant *and* timely to inform the CRM’s implementation of ERM at Watercare; even if he did extract some benefit from the collaboration in 2007 and 2008. Rather, because the thesis is the product of a lengthy (3-year) and detached period of sense-making, reflection, and critical writing, I can only claim that its contributions inform CRO practices in a more general sense.

I think this highlights an important point about the relevance versus rigour of transdisciplinary research (Huff and Huff 2001; Shani et al. 2008; Starkey and Madan 2001; Tranfield and Starkey 1998). As has already been noted, transdisciplinary research is defined in part by its attention to problems framed by stakeholders outside academia. Consequently, transdisciplinary research is also characterised by (i) complexity – the need to accommodate and integrate different disciplinary perspectives in order to grasp the complexity of the problem; and (ii) time sensitivity – the need to formulate workable solutions in a time frame that suits “real” world stakeholders. But these two features of transdisciplinary (and applied) problems are not necessarily congruent. Depending on the complexity and time sensitivity of the problem, and the available resources, it may not be possible to satisfy both academic standards of rigour and practical standards of timeliness for the delivery of solutions. Where this is the case, one must be sacrificed for the other. This author sacrificed timeliness in favour of rigour. As a result, the practical relevance of the findings from this research must be judged through their subsequent application in other contexts.

## **Judging the account: warrantability**

Watson (1997) points out that in addition to a practitioner readership, management research also has an academic readership. He suggests that the warrantability of an account is therefore also a question of whether the “conceptualizing, theory building and application of concepts to research observations is internally consistent and plausible in terms of current methodological debates” (1997, p 7). This chapter has attempted to locate the basis of this inquiry in the current resurgence of pragmatist methodology. It is, of course, up to the reader to judge the plausibility of my account, the cogency and consistency of my arguments, and whether I have achieved a successful marriage of context and theory with respect to the overarching project of the thesis. In forming these judgements, however, the reader should bear in mind that the thesis is not a straight-forward representationalist account of the CRM in action. This is explained further in the following discussion.

The presentation of the thesis is organised around two key themes (see Figure 2.3, overpage). The main thread is provided by the CRM’s paradigm shift, reflected in the order and empirical foci of the chapters (especially chapters 3, 4, and 5). There were three key parts to that shift, which together constitute the empirical narrative: the CRM’s initial vision for what ERM should look like at Watercare and his focus on “risk data” (chapter 3); the various interventions that the CRM pursued toward implementing that vision (chapter 4); and the various problems and dilemmas which resulted from those interventions and which ultimately caused the CRM to rethink his approach (chapter 5). The presentation of the narrative through chapters 3, 4, and 5 essentially maintains the temporal order of events. This was convenient for structuring purposes and, in combination with the thematic structure of the discussion (see below), provides the reader with a path through the conceptual territory of the thesis.

The narrative constitutes a claim as to what happened; that is, what the CRM did, what transpired in various times and places in which the CRM intervened in the organisation, and what the outcomes of those interventions were. The accounts of those moments are reconstructed from the audio transcripts (see the Dialogues), from observation notes that I made, from various sources of documentary evidence (e.g. emails, reports), and from memory. Where necessary the accounts are grounded in the broader social, organisational, and institutional context in which the moments took place so as to explain the particularities of what happened (for instance, in Chapter 4 the rejection by Watercare’s senior managers of the CRM’s assumption that the company’s Statement of Corporate Intent was a key strategic document is explained by placing the managers’ views in the broader context of the primary

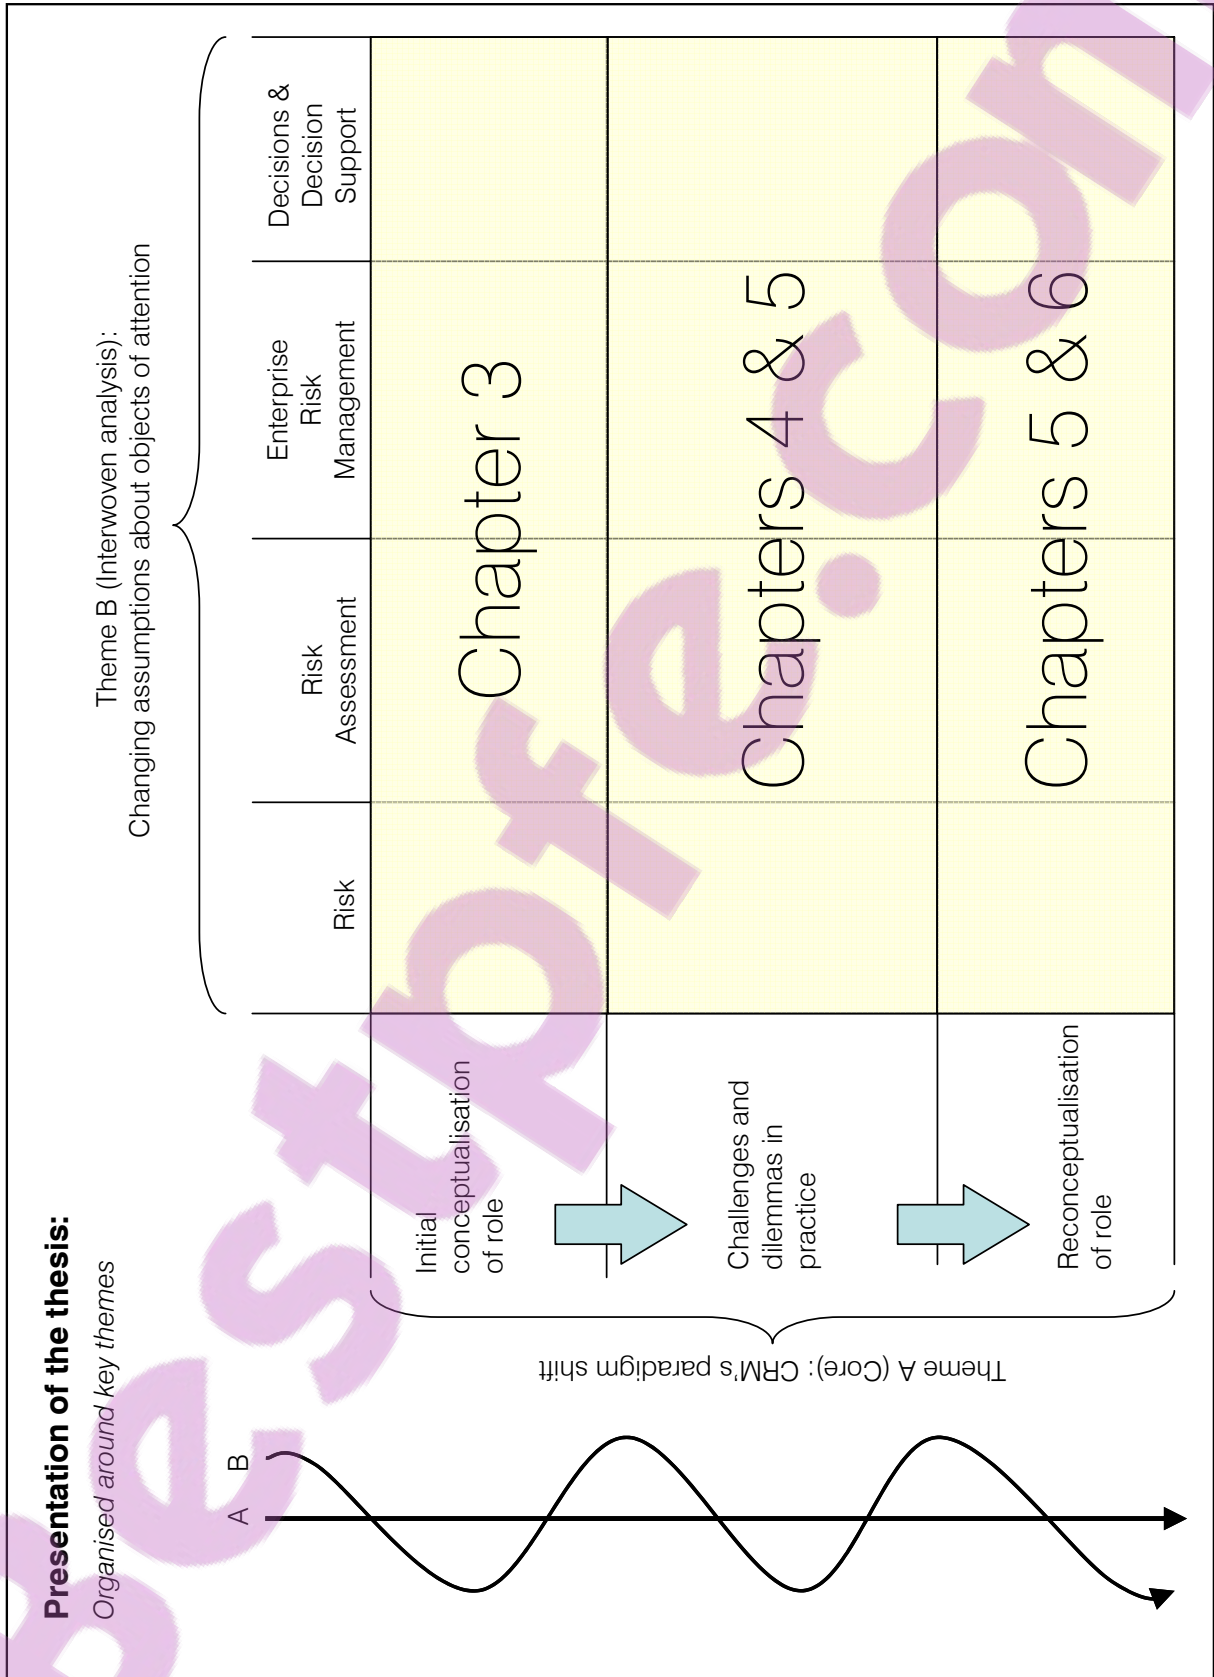


Figure 2.3. Presentation of the thesis

forces influencing the strategic direction of water and wastewater infrastructure development in Watercare's operating environment). In this regard I have tried to provide as full accounts as possible such the reader may see what I saw and thus be able to make up his or her own mind as to the significance and relevance of each moment to the arguments of the thesis. The reader may have to refer to the Dialogues and/or Appendices in order to access these accounts in full because priority has been given to presenting the analysis of them in the main body of the thesis.

The narrative does not, however, claim to describe the totality of the CRM's experience. Rather, it presents certain key moments selected from the CRM's broader programme of actions for their relevance and significance with respect to his paradigm shift. That much is left out of the account is unavoidable – one cannot be everywhere nor describe everything – but also raises the possibility that the account, and hence the arguments of the thesis might change in important ways for the inclusion of other moments. This is indeed a possibility. However, my selection and arrangement of the moments which make up the empirical narrative was not arbitrary or random, but rather was oriented by and in relation to the theoretical side of my project. Also, there is the fact that I was there, at Watercare, shadowing the CRM for 15 months. In that period of time, and as a result of my physical proximity to the sites of the CRM's performance, I acquired a pretty comprehensive understanding of his universe. The reader will thus have to accept that if there was something else of significance and relevance to the question of the CRM's paradigm shift then I would have included it.

The close and co-productive nature of my engagement with the CRM might be seen as problematic as far as the question of academic distance is concerned. As an actor simultaneously inside yet outside the organisation, there was a certain precariousness to my position since I was dependent on the CRM's (and Watercare's) continuing good will toward my project. This raises the question of whether I perhaps allowed my interests to be subjugated to those of the CRM or Watercare in order to retain my privileged access. As far as I am aware I have not done so either implicitly or explicitly. I managed this risk through the usual ethical process of declaring my interests to participants and requesting their permission, on condition of anonymity, to observe and record proceedings. All participants were given the opportunity to review transcripts and retract statements that they had made. In this regard, all of the project participants, and the CRM most of all, were generous with their time and access. I was never pressured by the CRM or anyone else at Watercare, nor did I at any time feel an obligation, to disregard or suppress information to which I had access that was relevant to my project.

There is a flipside to the above question: did I perhaps unduly influence the CRM's

thinking about his role? That is, by involving myself in his project, did I create a self-fulfilling prophecy? Certainly, to the extent that the CRM and I openly shared ideas and collaborated on various problems, we did influence each others' thinking. However, as I have already mentioned, we were not working on a common project, but rather two distinct but complementary projects, convergent around a common interest. Despite the imbalance of power just mentioned, neither of us expressed, implicitly or explicitly, a political interest to influence the outcomes of the others' project. In this regard, I believe our collaboration was objectively productive. I interjected information and ideas that I thought might be relevant or potentially insightful with respect to the CRM's project, and, vice versa, the CRM proffered ideas and reflections that he thought I might find interesting with respect to my project. But in neither case was there any direct influence. What I suggested to the CRM was subsequently subject to his reflections, and what he said to me was subject to mine. In true pragmatic tradition, we made use of ideas and theories to the extent that we thought them useful with respect to furthering our respective cognitive and practical endeavours. The lingering question of whether, this independence of mind notwithstanding, the CRM's direction might have been quite different had I not been involved in the project is irrelevant. First, I do not know and it is pointless to speculate on what the outcome might have been otherwise. Second, even if things would have been different, this does not make what did happen unworthy of study, nor render the insights that the thesis draws any less relevant to understanding the role of the CRO.

The second theme around which the presentation of the thesis is organised is the evolution of the conceptual basis of the CRM's decision support rationale. That is, the CRM's initial conceptualisation of his rationale, his questioning of that rationale in response to certain resistances and dilemmas, and his subsequent reconceptualisation of it are described and explained in terms of evolving assumptions about the key objects of attention; namely:

- **Risk:** What is Risk? How do we know Risk?
- **Risk Assessment:** What is Risk Assessment? What is its purpose? What are the features and characteristics of Risk Assessment?
- **ERM frameworks:** What is the purpose of such frameworks? How are they to be interpreted?
- **Decision support role of the CRO:** What is the goal of decision support? How is this goal to be achieved?

The various assumptions which characterised the three phases of the CRM's paradigm

shift are set out in Table 2.2 (p 68), as derived from the discussion sections of the thesis chapters. The reader should not, however, take the statements in Table 2.2 to be claims as to particular beliefs or assumptions that the CRM must have held, whether implicitly or explicitly. This sounds contradictory, but the point is as follows. Those statements and their thematic arrangement in Table 2.2 are the product of my analysis of the CRM's paradigm shift, but whether, when, and to what degree the CRM may actually have ascribed to those assumptions is not the claim that the thesis is making. Rather, the principal claim that the thesis makes is that, at least as far as supporting organisational decision making is concerned, there are different legitimate ways of conceptualising Enterprise Risk Management, of looking at the key objects of a Chief Risk Officer's attention, and thus of interpreting the role of the CRO. Those different ways of "seeing" imply different possibilities for acting, which may be more or less effective with respect to improving the quality of decision making in an organisation. It is argued that the perspective constituted by the right hand column of Table 2.2 is more likely to be workable as a basis for productive intervention by CROs than the other perspectives. This claim is not made about or in respect of Watercare's CRM, but about the role of CRO in general.

The methodological difficulty that the reader may perceive when deciding on the warrantability of this claim is that it is simultaneously the product of the interpretation and analysis of the CRM's actual performance, and of my own normative theorising about how the decision support function of the role of CRO should be approached. In other words, if the descriptive and normative aspects of the project were not independent, then how can either be trusted? My defense against any such challenge is the pragmatic, hermeneutic methodology described in this chapter. That is, the analysis of the CRM's performance was not biased by a preconceived normative ideal of the CRO role, and nor was the normative theorising of the role conducted with the aim of satisfying preconceived understandings of its practice. Rather both evolved together in a co-constitutive, dialogical fashion:

- The analysis of the CRM's performance was informed, guided, and challenged by evolving theoretical understanding of the role in a general sense;
- The theorising of the decision support function of the role was informed, guided, and challenged by an evolving analysis of the CRM's performance;
- The movement between the two sides of the project was iterative and reflective (see Figure 2.2, page 60);
- The questions and boundary constraints which guided the various stands of the



inquiry and the analytical framework around which the thesis is structured were not pre-given, but rather evolved through this interplay;

The above points also stand as a defense against the challenge that the CRM's status as a relative novice detracts from the value of his subsequent learning experience, and, consequently, from the contributions of this thesis. That is, while lessons learned from failure are often more important than those learned from success, what counts as an important lesson for one actor does not necessarily make a valuable contribution to the broader knowledge base for CROs more generally. Consequently, since the difficulties and confusion experienced by the CRM (and myself) were the product of his (my) inexperience and lack of knowledge with respect to the practical implementation of ERM, the lessons he learned (and hence the contributions of this thesis) are neither novel nor interesting when viewed from a position of such experience and knowledge. That the CRM was "learning the ropes" of ERM implementation was methodologically both beneficial and detrimental. On the down-side, it came with a significant increase in uncertainty and confusion (i.e. always feeling "muddled"), which translated into an increase in my workload with respect to sorting out what was relevant and significant. On the up-side, however, it also meant that the CRM moved slowly and spent a lot of time reflecting on whether he was going about things the right way. As I have already mentioned, the CRM's predisposition to reflection provided me important opportunities to access his thinking. As to the value of the CRM's journey, the thesis argues that the various problems, issues, and dilemmas that he encountered were not just peculiar to his performance, but rather are symptomatic of conceptions of ERM and the CRO role more generally.

*Table 2.2. Evolving assumptions about key objects of attention*

<b>Key objects of attention</b>	<b>Initial conceptualisation</b>	<b>Transition</b>	<b>Reconceptualisation</b>
Risk (What is Risk? How do we know Risk?)	Risk is something real: the quantifiable uncertainty (i.e. probability) of something undesirable happening; can be represented as Likelihood x Consequence.	risk = imaginable possibilities for the future, which may or may not be well known; in a broad sense risk is merely uncertainty about the future.  Risk = something both real and subjective; a specific form of knowledge about the future; a product of inquiry.	Sources of uncertainty in knowledge in a primarily epistemic sense. Actors may have more or less specific perceptions of these sources of uncertainty.
Risk Assessment (What is Risk Assessment? What is its purpose? What are the features and characteristics of Risk Assessment?)	A formal methodology or procedure to assess Risk, characterised by explicitation, specificity, and quantification.	risk assessment = the various ways in which people pay attention to uncertainty  Risk Assessment = a model for a rigorous (i.e. scientific) methodology for assessing and calculating Risk.	A methodology conceptualised as a convened function within organisational planning cycles which seeks to facilitate awareness, understanding, and prioritisation of uncertainties of attention based on an agreed vision of the world
ERM Frameworks (What is the purpose of such frameworks? How are they to be interpreted?)	Frameworks describe a generic process for the rational management of systems and processes; to be created or implemented within the organisation and integrated with extant business processes.	risk management = the various ways in which people pay attention and respond to uncertainty; in a broad sense is synonymous with management.  Risk Management = a generic model of good management; a quality standard for evaluating risk management.	A generic process model which provides a lens through which to view and interpret organisational processes in terms of Threats (negative uncertainties), Opportunities (positive uncertainties), and Controls (responses).
Decision support role of the CRO (What is the goal of decision support? How is this goal to be achieved?)	Good decisions should be based on good data. Therefore the goal of decision support is to produce good "risk data", which can be achieved by formal assessment and calculation of Risk.	The objectivity of information (i.e. good "risk data") is still the normative goal, but practical questions arise: What is the value of better data? What if good "risk data" is too costly to produce? How much detail is necessary for good decision making?	Good decisions are reasonable, but decisions are also personal. Therefore facilitate reflection about distinctions which constitute decision frames in order to reveal potential uncertainties.

## On how I became a transdisciplinary

As I already noted in Chapter 1, neither the subject nor the methodology of this study is typical of research normally conducted in the field of Engineering. A relevant question then is: how does an engineer come to perform such work? The answer to this question lies with the background to the research, which left me in a certain place at a certain time, with a certain set of interests, motivations, and dispositions, and certain connections to other actors which enabled me to recognise and take advantage of the CRM's project as an opportunity for further inquiry, and ultimately to be successful in that inquiry. There were three parts to that background, which are described in detail in Appendix II:

- The original project on which I started by PhD research with Watercare;
- The practical work that I performed on various process-improvement projects at Watercare during the period of the initial research (2004 – 2006); and
- The work of framing a new inquiry following the termination of the original research project.

The significant common feature of these early experiences was that they forced me to encounter and cross the boundaries of my engineering knowledge and experience. In each case I had to reach beyond the engineering knowledge base to develop insights and understandings.

In particular, my experiences with attempting to operationalise “sustainability” during the original research project foreshadowed in a significant way the later experiences of Watercare's Corporate Risk Manager with operationalising Enterprise Risk Management. The problems that I encountered were at root due to the limitations of the initial project approach for understanding and intervening in social systems to “engineer” certain outcomes. That initial approach was typical of the world-view and functional systems methodology in Engineering (the motivations, objectives, and approach of the original project are briefly described in Appendix II). The project methodology was, ironically, an attempt to specify a rational comprehensive method of inquiry to determine, scientifically, what the conditions for the “sustainable” state might be for water and wastewater infrastructure in the Auckland Region, which was intended to reveal the strategic changes needed to move development of the infrastructure in the “right” direction.

Within the paradigm of the project the concept of “sustainability” was well understood and the objectives and characteristics of a “sustainable system” were defined at least at a

conceptual level. In the terminology of Chapter 6, the research problem was assumed to be of a planning or operational nature. The objectives were known, the methodology was already defined, and all that remained was implement the methodology to determine the optimum solution. But during the course of the project it became increasingly apparent that the objectives of the project were at odds with the interests of the social system in which I was conducting the research. In essence, I encountered a very different notion of “sustainable infrastructure” from what had been assumed as the basis for my project. Eventually I realised that there were multiple legitimate ways of interpreting the “sustainability” concept in any social system, that the paradigmatic perspective from which my project had been defined was just one possible lens, and that it was not necessarily commensurate with the other competing paradigms which were encountered. This realisation forced me into the realm of strategic uncertainty (in the terms of Chapter 6) and thus in search of new paradigms and new approaches to cope with my situation. Ultimately, I realised that the project of operationalising sustainability required both new objectives and a new methodology which could acknowledge and accommodate that paradigm plurality.

Although I was unable to pursue the original research project to a conclusion, this early work was, in hindsight, vitally important in preparing me to follow and understand the later performance of Watercare’s Corporate Risk Manager (see Table 2.3):

- My background and training as an engineer provided me with a grounding in the engineering world-view and approach, as well as a sensitivity to the tendency of engineers to focus on calculative objectivity.
- The original research project provided me with the knowledge and experience applicable to conducting a messy sense-making inquiry of a complex, transdisciplinary role.
- From the original research project I also learned that there are multiple ways of interpreting (“seeing”) the world, and that different ways of “seeing” imply different possibilities for action (e.g. Sustainability as ...)
- The practical work provided me a critical sensitivity to the nature of the CRO role, the approach that Watercare’s CRM took to that role, and the issues ultimately encountered.
- When reframing the inquiry I sought to understand and theorise how the problems of water supply and sanitation were framed and solved through processes at multiple scales within the Watercare system. This work thus constituted my initial

interpretation of Watercare as a problem solving system, and an early attempt to blend theory and empirical insight via a pragmatic hermeneutic approach.

*Table 2.3. Prior experiences prepared me to follow Watercare's Corporate Risk Manager*

<b>Experience gained from the original research project (2004-2006)</b>	<b>Experienced gained from practical work at Watercare (2004 – 2006)</b>
<p>Experience, and perhaps a certain degree of comfort, with exploratory sense-making, especially looking across the discourse of multiple disciplines, merging theoretical insights from different disciplines with empirical insights from the field, and tracing (thinking through) connections, relationships, and contradictions to construct transdisciplinary understandings (models) of systems.</p> <p>Experience with the failure of a functionalist engineering methodology for designing interventions in social systems.</p> <p>Exposure to and understanding of world-views and approaches from the social sciences (i.e. the constructivist perspective and interpretive and post-structuralist methodologies); and especially the relevance of these with respect to intervening in social settings to change human behaviour.</p> <p>Experience working with a pluralistic and contestable concept which can be interpreted in a variety of ways across multiple disciplines.</p> <p>Experience with the problem of inquiring into and making sense of an organisation to establish “what matters” in practice.</p> <p>An interest in understanding how to change organisational behaviour to achieve different outcomes, which was carried across into the reframing of the new inquiry.</p>	<p>Exposure to and experience with performing some of the tasks of a Chief Risk Officer (e.g. investigating and understanding organisational structure and processes, and designing a risk framework).</p> <p>Sensitivity to the approach taken by Watercare's CRM (which was strikingly similar to the approach I brought to the projects I was involved with).</p> <p>Sensitivity to the issues encountered in that work (especially the question of the value of conceptual frameworks).</p>



# Chapter 3

## The CRM's vision for what ERM should look like

"I found myself concentrating on what... I wanted the risk management function to be able to deliver. I guess I focussed on the concept that risk management is supposed to support decision making and being an engineer I went to that thing that decisions will be based on data. So how can we provide data in a format and a level of detail that will enhance the way that we currently make decisions?"

*Watercare's Corporate Risk Manager explains his decision support rationale  
(Dialogue 1, p 239)*

Chapter 3 is the first of three chapters which examine the experiences of Watercare's new Corporate Risk Manager as he sought to implement Enterprise Risk Management within the firm. The purpose of the chapter is to reveal and characterise the CRM's initial approach to his role, and to justify that approach as reasonable despite his status as a novice Chief Risk Officer. Empirically Chapter 3 lays the groundwork for the analyses in Chapters 4 and 5 by framing up what the CRM wanted to achieve and how he intended to do it (i.e. specific proposals). Analytically, Chapter 3 lays the ground work for the discussions in Chapters 5 and 6 by making explicit (i) the underlying rationale which motivated the CRM's initial approach, and (ii) the corresponding normative agenda of ERM as a programme for the (re)construction of rational management in organisations.

The first part of the chapter describes the CRM's concerns with the quality of Watercare's "risk data" and his "vision" for what ERM should look like as it was expressed in his plan for developing the company's risk management capabilities. The CRM's vision is compared

against the distinguishing features of normative ERM models at the international level and against the calculative ideals expressed in engineering guidance for risk assessment. It is argued that the CRM's vision for developing Watercare's Risk Management function was consistent with both. The discussion section frames ERM as a normative programme for the (re)construction of rational decision making in organisations, and positions the CRM's approach as a specific interpretation of that agenda involving a particular focus on the objectivity of information. The CRM's approach is characterised, on that basis, as consistent with the "strategic controller" archetype of the CRO role described by Mikes (2010).

That the CRM initially approached his role in a certain way was a function of his inexperience with the Chief Risk Officer role, and of his background and training as an engineer. When he took up the position at Watercare, the CRM had:

- A professional background and high level education (PhD) as an engineer;
- Professional experience as a property risk assessor for an insurance firm;
- No prior senior management experience, and no prior experience with implementing ERM in practice;
- Limited conceptual knowledge of ERM gleaned from exposure to international ERM literature (e.g. standards, frameworks, and text books; especially Ward 2005), and from attendance at an ERM short course prior to taking up his position at Watercare.

Chapters 1 & 2 already laid some groundwork against the criticism that the CRM's novice status might somehow detract from the value of his journey (and the findings of this thesis). This chapter continues that work by demonstrating that the CRM's approach was in fact quite reasonable. That is, despite his relative naivety the CRM was not wildly misguided when he formulated his plan for developing Watercare's risk management capabilities. On the contrary, the CRM approached the role in a rigorous fashion, and his proposals were grounded within a coherent logic for the role of Risk Management within the organisation.

The CRM's evaluation of Watercare's risk management capabilities in the first part of this chapter contrasts quite starkly with the company's public statements (see Appendix I), and with perceptions of leading practice held at senior levels of the organisation, as well as by external auditors. It is not my intention to criticise or praise Watercare. My concern is with the CRM's journey and what that journey reveals about the practice of organising ERM. That the CRM initially perceived Watercare's risk management capabilities to be problematic in certain respects is important to that journey, and his views at that time serve as a point of entry into a certain paradigmatic approach to the Chief Risk Officer role.



## The CRM's proposals for developing ERM at Watercare

Upon joining the company in February 2007, the new Corporate Risk Manager immediately undertook to familiarise himself with Watercare's existing risk management capabilities by talking to management and staff and reviewing key risk management tools such as the corporate and project risk registers. The CRM also devised an intranet-based questionnaire in which he invited all company employees to give their frank (and anonymous) opinions on various aspects of the company's risk management practices. A formal audit of Watercare's risk management function by an external management consultancy in early 2007 provided a further independent perspective for the CRM to consider. The CRM identified areas of concern around the poor quality of what he called "risk data", the subjectivity of Watercare's existing risk framework, problems with risk communication within the company, and the apparent lack of maturity of Watercare's risk management function.

### Concerns about Watercare's "risk data"

The CRM quickly developed an overriding concern with what he perceived to be the poor quality of the company's "risk data". As he used it, the term "risk data" referred to the informational outputs of risk assessments performed by employees, which were formally recorded in the company's corporate risk register as statements of risk in a standardised format (i.e. cause, consequence, likelihood, severity, controls). In the CRM's opinion the data in that register exhibited a number of problematic symptoms. For brevity those symptoms are summarised in Table 3.1; a fuller description is given in Appendix IV. The CRM's concerns paralleled my own initial concerns about Watercare's "risk data" (see Appendix II).

*Table 3.1. CRM's concerns about the quality of "risk data" in Watercare's corporate risk register*

Symptom	Description
Poor risk descriptions	The CRM and the external auditors identified that the risk descriptions and assessments in the corporate risk register were often ambiguous, incomplete, overlapping, or simply too brief (Clement 2007e; PriceWaterhouseCoopers 2007). Entries in the register often described only the cause or the consequence of risks without clearly linking the two together, or lumped together multiple causal factors, each with different likelihoods, potentially different consequences, and different controls, in a single risk description.
Lack of transparency	The ambiguity and subjectivity of risk assessments was further compounded by the fact that the risk register was not transparent as to the quality of the

*Table 3.1. CRM's concerns about the quality of "risk data" in Watercare's corporate risk register*

Symptom	Description
	assessment behind each risk entry (Clement 2007e). The entries in the risk register did not record whether the final risk score was the result of an exhaustive, quantitative analysis, or simply a quick 'back of an envelope' assessment of someone's gut feeling.
Meaningless Enterprise Risk Profiles	The format for monthly risk reporting to the Chief Executive and the Board was to present the total number of risks in each business group and their distribution between five risk classes. But those representations were mathematically meaningless and therefore problematic as devices for communicating risk exposure: "...those numbers mean nothing, it's like throwing a handful of darts at a dart board and saying there's your profile" (CRM commenting on the crudeness of the enterprise risk profiles, 15 June 2007).
Lack of structure	The corporate register contained an overwhelming number of asset-focussed risks, typically each risk being defined as the loss or failure of a specific asset or group of assets (Clement 2007e). This meant that the enterprise risk profile was skewed and did not adequately reflect other important activities (and risks) across the company, and, since the entire body of corporate risks (~800) was contained within a single undifferentiated layer, there was no way to assess interdependencies between risks, or to generate meaningful cumulative risk profiles from the existing data set (Clement 2007e).

### Concerns about the existing risk framework (2003)

The CRM believed that Watercare's existing risk framework was the primary cause of the quality problems with the company's "risk data". At that time the existing framework had been in use within the company for nearly a decade, the last revision being in 2003 (see Appendix I for an overview of the content of the framework and a brief history of its development). The CRM perceived two problems with the framework (Clement 2007e): first, it was not properly aligned with the corporate objectives or with the real performance standards with which the company was required to comply; and, second, the terminology of the framework was arbitrary and vague.

The existing framework defined potential adverse impacts across five assumed performance categories which did not accurately reflect the strategic objectives of the firm. Unfortunately, the rationale behind the structure and content of the 2003 framework is not known in detail. The original consultant's report (Cooper et al. 1998) did not describe the particular process by which the categories, criteria, and descriptors in the framework were decided upon. One manager who had been involved in that process explained that the framework had been constructed by aggregating the objectives specified under the

Sustainability Policies in Watercare's 2003/04 Statement of Corporate Intent (SCI), but a comparison of the two documents revealed only a loose correspondence. The risk framework had not been updated despite changes to the SCI since 2003.

The evaluation criteria in the 2003 framework were also not consistently aligned with the real performance standards with which the company was required to comply, and against which the actual performance of the company was measured. This meant that the language of the framework did not reflect the specific terminology used by staff in the company's various operational contexts. Indeed, several of the criteria in the framework defined risk in terms of the consequential outcomes which could potentially arise from lapses in Watercare's performance (e.g. "environmental harm", "community outrage", "legal challenges", "disease"). Thus, strictly speaking, an evaluation of risk under the 2003 framework implied a double assessment: first, as to the magnitude and likelihood of a failure of Watercare's systems, and second, as to the magnitude and likelihood of any consequential outcomes potentially attributable to that failure. This amounted to judging how much trouble (legal or otherwise) the company would get into as a result of a breach of its performance standards, and such a judgement would always involve considerably more subjectivity than the primary calculation of the magnitude of the performance breach itself. Because of the way the framework was defined, this judgement was placed in the hands of the individual(s) performing the risk evaluation.

Finally, the consequence severity descriptors in the existing framework were, for the most part, qualitatively expressed in fairly general and vague terms. For instance, the framework made liberal use of terms such "negligible", "minor", "serious", "major", "significant", and "critical" without providing any guidance on what those terms actually meant in practice.

## **Concerns about risk communication**

In identifying the above problems with the existing risk framework, the CRM was cognisant of the fact that the framework was a core element of a boundary infrastructure (Bowker and Star 1999) which moderated communication between different communities of practice (Brown and Duguid 1991) within the company. Although the CRM did not use such abstract terms, it was clear from very early on that he perceived this as an important function, when, for instance, he noted how the development of the risk register had been driven by Operations staff as a mechanism to make their "worries" visible to management (notes from interview with CRM, May 2007). In effect, the framework provided a common language with which management and staff could communicate about the company's risks.

The CRM felt that the existing framework was problematic in this regard because it arguably only represented the organisation from a strategic or top-down perspective:

...the current framework... I think was really an arbitrary categorisation of the strategic objectives, and they've gone like this [makes dumping motion with hand]. Now that's fine, that's a representation of sorts, but there's clearly a disconnect between the way the framework represents the business and the way they think about and analyse it at the coal face. (Dialogue 9, p 305)

When he made the “dumping” motion with his hand in Dialogue 9, the CRM was implying that the designer(s) of Watercare’s existing framework had dumped or imposed on the organisation an essentially arbitrary categorisation of the corporate objectives. The CRM saw this as problematic because if the framework lopsidedly reflected the objectives and terminology of only one community of actors in the organisation, then others would find it frustrating and meaningless to use; which is what happened at Watercare. In order for Watercare’s staff to communicate their operational knowledge to management, they had to translate that specific knowledge into the unfamiliar and arguably arbitrary terms of the risk framework. Not only was this translation individually subjective, but it was further compounded by the vagueness of the language with which the framework was defined, and by the fact that the framework required the assessor to imagine what *might* happen in the broader environment and community in the event that Watercare failed to perform to standard, and what others (e.g. stakeholders, customers, the public, regulators) *might* do about it. This internal subjectivity was a source of inconsistency in risk evaluations because different assessors would inevitably judge similar situations quite differently.

The CRM expressed concern that the resulting ambiguity and subjectivity in risk assessments would lead to problems in understanding and communicating risk information: “Ambiguity in the risk description makes communications highly inefficient and often ineffective. In particular poor risk descriptions make it difficult to gauge what exactly is covered (and not covered) by the risk description, what exactly is being discussed or evaluated, and what the most effective method would be for risk control” (Clement 2008c, p 10). Indeed, the CRM found that this was exactly what happened in practice:

A comment from a single staff member summarised my suspicions nicely. We had just been in a workshop with six senior managers where we were reviewing the evaluation of various risks. The first hour of the workshop had been spent debating one single risk, both what we were talking about and then how it should be scored (in terms of consequence and likelihood measurements). After the workshop one of the senior managers said to me – “that’s what we normally find happens, it takes us about an hour or so to start thinking along the right lines, then we get into the swing of things and tend to whip through the rest of them”. In actual fact, what

was happening was that after an hour of debate the workshop attendees gave up trying to resolve and agree what they were talking about, and subconsciously accepted that the risk description contained enormous ambiguity and that in order to finish the workshop they needed to stop trying to think about the detail. To me, it reinforced just how much valuable time could be wasted. We were consuming a significant amount of senior employee time, and yet they were glossing over the detail, which is generally where the devil resides. (excerpt from Clement 2008b, presentation to ERM Conference, April 2008)

The CRM further identified that too much subjectivity in the risk framework was leaving the risk assessment process open to manipulation (Clement 2007e); a view which was also expressed by various staff members, both directly to the CRM (see Dialogue 5) and in responses to the staff survey (Clement 2007c). There was clearly motivation for staff to exaggerate the expected risk reduction in capital expenditure requests, not only to overcome the limitations of the RMS scoring method, but also because staff generally perceived that minor risk reductions would not attract investment (Clement 2007c). Indeed, since bigger budgets often equate with greater responsibility and perceived importance within the organisational milieu it is understandable that actors would tend to inflate the significance of risk outcomes to secure operational or capital funds for their area of the business. It was, somewhat perversely, a matter of pride amongst certain actors within Watercare to be able to claim that their risks were bigger than anyone else's (this can be seen to some extent in Dialogue 4) – a feature of the corporate culture which can be traced to the strong emphasis on risk as the basis for decisions over the allocation of capital and operational funds. Ironically, however, it was also widely understood amongst staff that it was culturally unacceptable to present Class 5 risks to decision makers. This created pressure to ensure that the risk score for the status quo was high enough but not too high, and also added to the difficulty of demonstrating significant risk reduction from the proposed expenditure.

The significance of the above issues was revealed in a meeting between the CRM and several line managers from Watercare's Operations group in August 2007 (see Dialogue 4). The CRM had called that meeting early in the development of his revised risk framework to address the question of how to represent the corporate objectives with concrete parameters that could be quantified using common business data. The discussion in that August meeting revealed a broad diversity of perceptions amongst members of different communities of practice within the organisation about "what mattered" with respect to risk outcomes (see the analysis of Dialogue 4 in Appendix VI, under the sub-section *Engaging with Stakeholders*). Since there was little in way of an objective point of reference in the existing framework, evaluations of the severity of risk consequences would inevitably be grounded in the perceptions of individual assessors; perceptions which could, and indeed did, involve

considerable subjectivity. This pointed up the considerable potential for inconsistency and mis-interpretation of “risk data” under the existing framework, and reinforced the perception that the existing risk framework was problematic as a moderator of risk communication.

### **The evidence from the CRM’s survey of staff perceptions**

The CRM developed a survey questionnaire to elicit staff perceptions of Watercare’s Risk Management function. The survey was hosted on the company’s intranet for one month (August to September 2007) and all employees were invited to participate. The results of the survey were kept anonymous. Approximately 28% of the employee population (97 participants) responded to the survey. Although the CRM was not formally trained in the practice of survey construction and analysis, the form and content of the survey and his analysis of it reflected a rigorous logic. The results of that survey can therefore be taken as an indicator of the state of Watercare’s “risk culture”.

Feedback from senior management indicated that there was a relatively strong culture of risk awareness at senior levels within the organisation, that roles and responsibilities with respect to risk management were clearly established and understood, and that the risk framework was valued by senior staff as an important mechanism for communicating their priorities (PriceWaterhouseCoopers 2007; notes from interview with CRM, 17 May 2007). The results of the CRM’s survey, however, revealed a picture of the company’s “risk culture” which was not consistent with perceptions at the senior management level. Key results from the CRM’s survey were as follows (Clement 2007c):

- Half of the survey respondents believed they had either no or only rudimentary knowledge of risk management concepts or theory, and over half had received either no training or only informal ‘on-the-job’ training in risk management. Nearly 40% of respondents indicated that they had not yet been exposed to the WSL risk management system in their roles.
- 60% of respondents claimed that they had never tried to use the Corporate Risk Register. The same percentage either did not know or could only speculate about the major risks facing their business units, and only one third admitted to a clear understanding of what risks/risk controls they were responsible for.
- Half of the respondents said they were hardly ever involved in risk management discussions, and of those who were involved in such discussions, two thirds found them confusing, vague, and open to misinterpretation.

- Nearly 70% of respondents indicated that the Capex template was not applicable to their roles. Of those respondents involved in preparing risk analyses for capital expenditure applications, three quarters believed the analyses were subjective and open to manipulation.
- 65% of respondents indicated that the Project Risk Register was not applicable to their roles. Of those respondents involved in project management, over half rated the company's risk framework and scoring method as unsuitable for evaluating project risks.
- Half of the respondents rated risk management as providing little or no value to their roles; only 13% believed risk management provided valuable insight.

The results of the CRM's survey contrasted somewhat starkly with the idea that Risk Management at Watercare was best practice, and that "risk data" was extensively used to support organisational decision making. To the contrary, the results of the survey seemed to indicate that risk management was not well integrated as a formal process or culture within the business, and, largely due to the data quality issues described above, was not convincingly delivering the added value expected of modern risk management functions (Clement 2007c).

### **A vision for better "risk data"**

The CRM's review of Watercare's risk management capabilities formed the basis for his three-year vision and development plan in which he proposed a number of initiatives to address the identified weaknesses (Clement 2007e, 2007f). The primary focus of that plan was on improving the quality of risk data by developing capabilities for the objective assessment, calculation, and representation of risk profiles, both within business group silos and integratively across the enterprise. It targeted four fundamental improvements:

- i) Development of an explicit risk hierarchy devolved from the corporate objectives to mirror the organisational and informational structure of the organisation;
- ii) Integrating the risk register with existing business systems and processes so that business data could feed into risk assessments;
- iii) Application of structured and systematic methods to promote comprehensive risk identification and complete and unambiguous risk description; and
- iv) A shift to quantitative techniques for calculating likelihood and consequence (i.e. risk

modelling) to enable the generation of meaningful aggregate risk profiles.

The CRM's proposals were outlined in two primary documents, his Framework Development Plan of March 2007 (Clement 2007c) and his Vision of May 2007 (Clement 2007d). Table IV.1 (Appendix IV, p 441) collates and summarises the specific proposals from each document under common functions. The following sections explain the key proposals.

### **A detailed risk framework and integrated risk register**

The Corporate Risk Manager felt that an integrated risk register linked to a rigorously defined risk framework would provide an efficient solution to the problem of collecting, aggregating, and reporting of 'risk data' within the company. The CRM expressed the vision as follows:

My initial thoughts focussed on having a data architecture within the risk register that made linkages between the risks in accordance with the structure that would be generated through fault tree analysis. I had envisaged a hierarchy of risks, derived from the principal objectives of the organisation at the top, with layers of subordinate risks at various levels of detail below. (excerpt from Clement 2008b, presentation to ERM Conference, April 2008).

The aspiration was that the architecture of the risk register should reflect the operational and informational structure of the organisation:

The data structure in the risk register should ideally reflect the hierarchy of authority in the company so that risk reporting can be made directly relevant to specific levels of management. The most effective method to structure risk data is to explicitly link each risk to a set of objectives and priorities set by the organisation's leadership and cascaded down through the management structure. (excerpt from Clement 2007f, p 8)

Each layer of the risk hierarchy would be mapped to corresponding management functions thus aligning the risk hierarchy with the management and control hierarchies in the organisation. The intention was that data generated from business processes would feed into the bottom of the register and be aggregated upwards to calculate meaningful risk profiles at higher levels. Managers would then be able to interrogate the risk register and see up-to-date profiles for the risks relevant to their position, at a level of detail appropriate to their degree of control (Clement 2007e; see also Dialogue 1). In this way, the risk register would become, in principle, dynamically integrated into the business as a functional management tool, rather than existing as a detached or "bolt-on" repository of static information (Clement 2008b).

The CRM's vision for an integrated risk register required the redefinition of Watercare's existing risk assessment framework. The CRM believed that the assessment framework



should be, as far as possible, an objective standard for comparing and evaluating the significance of risk outcomes right across the organisation (this being a pre-requisite for ERM). The CRM proposed that the framework would be reformulated from an explicit deconstruction of Watercare's Statement of Corporate Intent, since the objectives contained within that document provided, at least in principle, the common and unbiased basis on which to define risk consequences at an enterprise level (Clement 2007a). The CRM expected that the resulting risk hierarchy would more accurately represent the various operational contexts of the business, which would not only facilitate the comprehensive identification of risks across the business (i.e. the hierarchy itself would highlight where potential 'holes' existed in the 'risk data'), but would also provide assurance that the risk management function was supporting the corporate objectives.

In order to address the identified problems with the subjectivity of risk assessment and communication under the existing framework, and to support his vision of generating quantitative enterprise risk profiles (see below), the new risk framework would need to contain objective performance measures: "[t]he thing that stands out immediately is that if you want to quantify [those risks] then you need some measure of performance, so some measure of service delivery standards or requirements, and some measure of business process or management requirements." (Dialogue 8, p 299). The framework would also need to be comprehensive, "[i]t's about having an objective way of measuring things across an organisation, full breadth and full depth." (Dialogue 3, p 257). If the framework was too simple and generic then it would neither address the problems of subjectivity and ambiguity which plagued the existing (2003) framework, nor provide for the calculation of quantitative enterprise risk profiles.

To this end, the CRM believed that the consequence evaluation criteria in the framework should be drawn from the relevant performance standards under which the company operated, both external (i.e. statutory obligations and the customer contracts) and internal (i.e. the operational budget, the Asset Management Plan, and the Funding Plan). For instance, one of the relevant external standards for potable water quality was the New Zealand Drinking Water Standards, which specified a range of water quality related criteria. The CRM anticipated that it would be possible to collate or otherwise rationalise those criteria to produce appropriate indicators to represent risk impacts on water quality.

### **Systematic analytics and quantification**

The ability to calculate and represent meaningful risk profiles at different levels in the

proposed risk hierarchy was central to the CRM's vision for the development of the risk management function at Watercare (this was the main theme in Dialogue 1). Such a capability would require rigorous, systematic risk identification, clear distinction between causes and consequences in risk descriptions, quantification of both likelihood and consequences using common, consistent measures across the organisation, and probabilistic methods of risk aggregation (Clement 2007e, 2007f). To improve the probability that risks would be identified systematically and comprehensively, the CRM proposed to facilitate the adoption of recognised, structured methods for risk identification (Clement 2007e). Although he did not identify specific methods, the CRM was presumably referring to well established methods for risk identification and assessment described in any engineering risk management text book (e.g. Ayyub 2003). The CRM did, however, seek to address the tendency of staff to confuse causes and consequences when describing risks by introducing a structured "risk description template" (Clement 2008d). The template explained how to create clear and complete risk descriptions by linking consequences with specific causes. He also undertook to redesign the corporate and project risk registers to improve the collection of risk data (this work is described in Appendix VI and discussed in Chapter 4).

The CRM also envisaged that different types of models could potentially be utilised to quantify likelihoods and consequences for different types of risks, which would then enable the use of Monte Carlo-based simulation to quantify cumulative risk profiles. The ultimate vision was that such simulations would provide a more objective and robust basis for the comparison and justification of different maintenance and capital expenditure programmes. Watercare's reliability centred maintenance (RCM) programme served as both an inspiration and ultimate goal. By mid-2007 a Reliability Centred Maintenance programme had been in place at Watercare for three years, and a number of detailed models of the company's water and wastewater treatment plants had been developed. For the CRM, the logical next step was to extend the RCM programme to the rest of Watercare's physical infrastructure, i.e. to the water distribution and wastewater collection networks (see Dialogue 2). The target functionality was to be able to calculate aggregate risk profiles that were meaningful (or at least 'improved') reflections of the company's real risk environment (Clement 2007e, p 9), and which were sufficiently robust, in terms of the underlying method, to be presented to external stakeholders as explicit and transparent justification of Watercare's capital investment programme (Clement 2007f). This was the basis of the CRM's proposed Investment Uncertainty Evaluation and AMP Formulation tools (see Table IV.1, pp 441-443), the former seeking to quantify risk profiles on individual investment decisions (projects) and the latter seeking to quantify risk profiles on the company's entire capital investment portfolio, the

Asset Management Plan (Clement 2007e; 2007f; interview with CRM 20 July 2007; see also Dialogues 2 & 3). This kind of robust support for the AMP was something that the CRM felt would become increasingly important in the future (Clement 2007d).

## Comparing the CRM's proposals against normative models

Appendix V describes the key features of Enterprise Risk Management as it is depicted in prominent international frameworks and standards. The central component of ERM is the generic Risk Management process, which is envisaged as being applied across all business functions and integrated throughout all business processes. Enterprise Risk Management is distinguished from isolated performances of Risk Management activities, in specific times and places within an organisation, by three key features: (1) taking a cross-functional, enterprise-wide perspective of risk, (2) embedding risk management throughout the organisation and its processes, and (3) the explicit centering of the corporate objectives as the focal point for all encounters with risk (see Appendix V). The CRM's vision for developing ERM at Watercare reflected all three of these features (see Table 3.2, overpage).

Capability Maturity (CM) models for Enterprise Risk Management describe the general characteristics of ERM processes at different levels of maturity (see the summary of these in Appendix V). Initially Risk Management is either non-existent within the organisation, or performed only on an ad hoc basis. At intermediate levels Risk Management is a distinct, explicit, and formal process within the organisation, consistently applied. When the Risk Management capability is fully developed actors are said to proficiently and consistently perform the process, it is "fully integrated" as a natural part of all processes and roles in the organisation, and the process is formalised and explicit in its execution, and supported by specific tools and calculative infrastructure, and there is an explicit role for science and evidence in decision-making (MacGillivray et al. 2007a, 2007b; Pollard et al. 2004; Ward 2003, 2005). Ward, for instance, commented that:

...fully integrated RM would imply, among other things, that RM was applied in all organisational decisions, large or small, strategic or tactical, complex or simple. This corresponds to Hopkin's (2002) notion of 'aligned' and 'embedded' RM that supports strategic and project plans and operational procedures" (excerpt from Ward 2003, p 13).

Proponents of Capability Maturity models claim that they provide a normative guide for the development of ERM in organisations. Pollard et al. (2004, p 460), for example, note that the value of such models "is in identifying what measures are required for an organization to progress between levels of maturity in risk analysis and management", although qualified by

**Table 3.2.** Comparing the CRM's vision against normative ERM principles

Normative elements of ERM (Appendix V)	Elements of the CRM's vision
<p><b>Centering the corporate objectives:</b> ERM seeks to impose a common definition and basis for measuring risk so that risks assessed in one part of the organisation may be comparable with risks assessed in others. That common point of reference is the organisation's strategic objectives, and other such objectives which flow "from the strategy, cascading to entity business units, divisions, and processes" (COSO 2004, p 18). Centering the corporate objectives is essential to the ideal of the Enterprise Risk Profile (i.e. accounting for <i>all</i> identified sources of risk across the enterprise).</p>	<p><b>Centering the corporate objectives:</b> The CRM proposed to deconstruct the objectives in the Statement of Corporate Intent as the common and unbiased basis on which to define risk consequences at an enterprise level (Clement 2007a). He proposed that the existing one-size-fits-all framework should be replaced by three inter-related frameworks, for the Operations, Project, and Strategic management contexts, constituting a hierarchical system cascading from the strategic objectives to reflect the specific objectives, time scales, and magnitudes of risks in each context.</p>
<p><b>Creating the risk culture:</b> ERM envisages that risk management should be embedded throughout the organisation as a culture rather than as a "bolted on" bureaucratic function (Lam 2003; Layton and Fuchs 2007; Shenkir and Walker 2006; Ward 2005). The basic idea is that RM is more effective and efficient when it is carried out as an integral part of organisational processes, rather than as an added extra that people have to accommodate on top of their normal work. The ultimate vision is of a risk-aware culture where people are aware of and understand risk in the context of the organisation's objectives, where <i>all</i> decisions in the organisation involve the explicit consideration of risk, and where risk management procedures and tools are employed as a matter of course in day-to-day activities (COSO 2004; ISO 2008; Shenkir and Walker 2006).</p>	<p><b>Creating the risk culture:</b> The CRM proposed to introduce regular, formal training to maintain and develop the knowledge and competency of staff in the practice of risk assessment, particularly with respect to the rigour of risk identification, and the proper and complete description of risks. He also expected that the new risk framework, aligned with the functional and informational structure of the organisation, and defined in terms of the company's real performance standards, would facilitate the cultural embedding of risk management practices within the organisation. That is, the CRM anticipated that because the structure and data in the new risk register would more accurately reflect the various operational contexts of the organisation, it was more likely to be utilised as a valuable management tool.</p>
<p><b>Developing the cross-functional view:</b> ERM is holistic and cross-functional in contrast to the traditional treatment of risks in functional "silos" (Lam 2003; Power 2007; Ward 2005). The approach seeks to address interdependencies between risks at the enterprise level that would otherwise be missed by traditional approaches to risk management. The central concept is that of the Enterprise Risk Profile, conceived in abstract terms as an aggregate representation of the firm's risk universe (Cummings 2008), accounting for all types of risk exposures across different business silos, and for the inter-relationships between those exposures.</p>	<p><b>Developing the cross-functional view:</b> One of the CRM's primary concerns was that existing representations of risk, particularly at the enterprise level, were essentially meaningless due to the subjectivities involved in the risk assessment process and the improper manipulation of risk scores. He was also concerned that the existing risk register was significantly biased toward Operations (i.e. asset-failure) risks. Developing the capabilities to assess and evaluate risk consistently across the organisation, and to calculate mathematically meaningful, quantitative, aggregate risk profiles at the enterprise level were central objectives of the CRM's vision.</p>

**Table 3.3. Maturing Watercare’s Risk Management (RM) capabilities**

<b>RM Aspect</b>	<b>Current (perceived) and future (envisioned) states of capability maturity</b>
Policy and process definition	<p><u>Current:</u> Watercare’s RM capabilities remain relatively immature and siloed.</p> <p><u>Future:</u> Watercare will be a proficient practitioner of advanced ERM functions.</p>
Culture	<p><u>Current:</u> Risk culture well developed at senior levels of the organisation, but only limited development within main body of employees. Limited understanding of risk concepts beyond individual asset failures. Risk analysis integrated into capital planning processes, but RM otherwise treated as a “bolt-on” function.</p> <p><u>Future:</u> RM processes institutionalised as part of “the way things are done”; culture permeating full depth and breadth of the organisation (i.e. everyone is aware of the risks in their context and proficient in the performance of risk management processes).</p>
Risk assessment framework	<p><u>Current:</u> Defined but lack of structure and strategic focus; criteria arbitrary and vague.</p> <p><u>Future:</u> Fully defined as a detailed hierarchy cascading from strategic objectives; criteria defined in terms of actual performance standards to which organisation is accountable.</p>
Data quality	<p><u>Current:</u> Lack of structure in the risk register (single undifferentiated layer), poor quality data (ambiguous, subjective, incomplete, overlapping risk descriptions), and silo view (bias to asset operations risks). Most risk knowledge contained within heads of experienced individuals. Problematic representation of risk, particularly aggregate risk profiles.</p> <p><u>Future:</u> Risk register restructured to reflect the management hierarchy of the organisation and fully integrated into business processes as a key management tool and repository of reliable, up-to-date risk information. Risk properly represented, with aggregate risk profiles developed through Monte Carlo simulation.</p>
Risk analysis	<p><u>Current:</u> Sophisticated reliability modelling applied to water treatment plants, but otherwise relatively basic risk analysis capabilities.</p> <p><u>Future:</u> Fully developed analysis capabilities, including cross-functionally, with a range of qualitative, semi-quantitative, and quantitative tools available to suit purpose and context.</p>

recognition that resources are finite and that capabilities should match the needs of decision-makers. Table 3.3, above, illustrates the CRM’s intention to develop Watercare’s risk management capabilities toward a state of relative maturity and proficiency by characterising Watercare’s existing risk management capabilities, as perceived by the CRM, against Watercare’s future risk management capabilities, as envisioned by the CRM, in terms of key indicators of ERM capability maturity (cf. Tables V.2, p 451, and V.3, p 452)

The concerns that the Corporate Risk Manager developed about Watercare’s “risk data” and risk assessment framework were also strongly influenced by his training as an engineer,

which equipped him with a certain paradigmatic view of the world and a certain methodological approach to knowing and intervening in that world. For engineers, risk arises from the behaviour and interactions of real systems and processes, which must be understood through scientific analysis. The engineering approach to risk assessment is characterised by a functionalist systems methodology: an assumption that the real world is systematic and can be analysed in systems terms, systematic analytical methods which take into account sub-systems and components through the construction of systems models, and an emphasis on quantitative methods for assessing the failure probabilities and consequences of engineered systems (Auyang 2004; Ayyub 2003; Jackson 2000; see Appendix V, p 452-464). This perspective was highly compatible with the dominant notion of risk within the cultural milieu of the Watercare organisation, which was associated with the potential failure of physical assets within the water and wastewater infrastructure networks. While this may be seen as a form of bias from an ERM perspective, which demands an integrative view of risk encompassing all of the risk ‘silos’ across the enterprise, it was not unreasonable in light of the fact that developing, operating, and maintaining infrastructure assets was Watercare’s core business. As such, the company’s primary business activities were engineering activities, and the bulk of the company’s employees were engineers, engineering technicians, or scientists. Thus, at Watercare, engineering concepts of risk and engineering approaches to risk assessment predominated.

In this context, and in light of his background, it is not surprising that the CRM perceived the state of Watercare’s “risk data” as virtually the antithesis of “good” data in an engineering sense; being subjective, ambiguous, incomplete, and non-transparent, and, therefore, not a properly objective representation of Watercare’s true risk profile. It is also not surprising that the CRM subsequently rationalised Watercare’s poor “risk data” as a symptom of having a subjective risk framework, a lack of systematic and quantitative methods for risk assessment, and the absence of a properly structured risk register; or that his proposals for addressing the data quality issue centred around the development of frameworks, hierarchies, structures, and methods of calculation (see Table 3.4). Indeed, the core proposals of the CRM’s vision read like a text book on engineering risk assessment; see Appendix IV.

I can attest to the influence of the engineering perspective in this regard. Prior to following Watercare’s CRM I worked on various process improvement projects which sought to modify aspects of Watercare’s capital decision infrastructure, including the risk assessment framework. My reflections on my approach to those projects (see Appendix II) revealed themes which strongly paralleled the CRM’s concerns about the quality of Watercare’s “risk data”, the need to improve the calculation and specification of risk, and a

belief in the need for proper process specification, detailed frameworks and regulatory infrastructure. I believe that my training and background as an engineer strongly influenced my attention to those issues.

*Table 3.4. The CRM's desired improvements for "risk data" quality*

Development proposal	Improvement's sought
Classification frameworks	Comprehensiveness, clarity, consistency, resolution
Measurement and data capture	Accuracy, comprehensiveness, frequency, resolution
Procedures and methods	Rigour, analytical robustness
Calculative infrastructure	Mathematical accuracy, complexity (modelling complex phenomena)

### **The CRM's rationale: supporting and defending decisions with "good risk data"**

Watercare's CRM expected that, properly implemented, ERM should deliver two primary outcomes: (1) *Insight*: ERM would improve the company's analytical capabilities so as to provide "greater and clearer insight into complex or critical decisions" (Clement 2007f, p 4; this outcome reflects the decision support function of ERM); and (2) *Defensibility*: ERM would make decisions defensible by providing the organisation with the capability to explicitly demonstrate due process and a robust basis for decision making (this outcome reflects the governance and internal control function of ERM). The CRM believed that both outcomes should be achieved through the production of good "risk data"; a rationale which he expressed from very early on:

I found myself concentrating on what... I wanted the risk management function to be able to deliver. I guess I focussed on the concept that risk management is supposed to support decision making and being an engineer I went to that thing that decisions will be based on data. So how can we provide data in a format and a level of detail that will enhance the way that we currently make decisions? (Dialogue 1, p 239) .

### **Providing insight through "good risk data"**

The CRM's notion of "decision support" meant providing information to decision makers that would be accessible, transparent, timely, and at the right level of granularity (detail) for the

decision being made (notes from interview with CRM, 17 May 2007). The CRM expressed the vision as follows:

The strategic aim of the three year development programme is to change the internal perception of how risk management delivers value for Watercare... Improving the collation of information will focus on reducing ambiguity, improving data structure, delivering transparency (common understanding) and improving integration with business systems. As a package these advances will improve the efficiency of the risk management function for staff, to better facilitate uptake. Improving the risk-based support for decision taking will focus on developing the analytical capability of risk management tools. This will provide greater and clearer insight into complex or critical decisions to improve the quality of decision taking. (excerpt from Clement 2007f, p 4)

It could be suggested that this rationale conflates “risk data” with what actors might actually know or understand about a given set of risks. Certainly, data and knowledge are not the same thing (Bell 1999; Tsoukas and Vladimirou 2001), but I do not think that the CRM (nor myself) ever believed in a 1-to-1 correspondence between data and knowledge. Rather, we both assumed that “risk data” performed, or should perform an informative function, which the CRM expressed as providing insight. That is, the better the data or information at one’s disposal the more informed one should be about the decision at hand, and, all else being equal, a more informed actor should make better decisions. Good “risk data” would perform such an informative function as an objective representation of the systems, processes, and environments about which, and within which, decisions are made. It was through this mechanism that ERM would fulfil the *value-creation* claim expressed in the world-level ERM literature.

In this regard, the CRM’s concerns about the quality of Watercare’s “risk data”, and his perception that risk assessments were not widely used to support decision making in the operation context, constituted a strong imperative for change. On the one hand, risk was a key factor in all capital decisions, but the CRM was concerned that risk evaluations performed under the 2003 framework were producing risk profiles (i.e. “risk data”) bearing only an unreliable and inaccurate relation to the real probability of future impacts on the company’s performance (as defined by the actual performance standards). The framework did not objectively reflect the company’s actual corporate objectives and performance standards, and therefore implied a significant degree of subjectivity in the assessment of risks. There was evidence that risk evaluation and communication was being adversely affected by that subjectivity. The risk entries in Watercare’s risk register were evidently incomplete, ambiguous, and subjective, as well as being non-transparent as to their sources. And the graphical representations of Watercare’s risk profiles were mathematically



problematic. On the other hand, it appeared that risk assessments were not widely employed in the operational context, leading the CRM to conclude that decision making at Watercare was informed by either poor quality “risk data” or no “risk data” at all. This suggested problems with the quality of staff knowledge about the company’s risk exposures more generally; a view which was apparently supported by the results of the CRM’s staff survey. This constituted a strong imperative for the CRM to target his interventions toward improving the quality of the company’s “risk data”.

### **Defending decisions with “good risk data”**

The CRM’s expectation was that ERM should support organisational governance and internal control processes principally through the production of good “risk data”:

When I first got here, I recall using the term “defensibility” a lot. I was talking a lot about the risk data. I thought that the risk management system had a huge role to play in terms of providing the organisation with defensibility for the way it made its decisions. So the data should justify the decision making and back then I wrote that generating data was one of my fundamental roles, to justify or provide a basis for decision making. (Dialogue 14, p 331)

The reason we generate and analyse risk data is to provide objective justification for our decision making. (excerpt from Clement 2008b, presentation to ERM Conference, April 2008)

In this case “risk data” would constitute an objective record of the decision making process and its outcomes. Such a record is necessary not just for the ex post facto defence of decisions but, from the point of view of Internal Control, it is also essential for the ongoing monitoring of actual performance against objectives set during the decision making process (this includes monitoring the effectiveness of risk controls). It was through this mechanism that ERM would fulfil the *governance* claim expressed in the world-level ERM literature.

The CRM’s review of Watercare’s corporate risk register led him to conclude that the quality of Watercare’s “risk data” was problematic in this respect: i.e. the representations of risk contained within that database did not accurately or reliably reflect the company’s true risk exposure, or the range and importance of the various controls which the company employed to mitigate those risks. As such that data could not be reliably used to justify decision making to external stakeholders or as a baseline for internal audit or control. This constituted a strong imperative for the CRM to target his interventions toward improving the quality of the company’s “risk data”.

## **Discussion: ERM and the construction of rational management in organisations**

World-level standards and frameworks for ERM have been described as generic models or templates for the (re)construction of rational management in organisations (Power 2004, 2007). The core component of the ERM model is the Risk Management process, normatively defined as a multi-stage procedure for the rational management of systems, processes, and decisions, to be applied to and across the functional areas of the business, and integrated throughout the enterprise, top to bottom (COSO 2004). That procedure requires the explicit identification and assessment of uncertainties with respect to the performance of business systems, processes, and decisions. The implied value claim is that if the system and its performance objectives are comprehensively defined, and if risks are comprehensively identified and rigorously assessed, then threats are more likely to be mitigated and opportunities are more likely to be realised, leading to improved and more reliable performance, compared with implicit, reactive, and unplanned approaches to management. Thompson (1986) cleverly illustrated this claim with a story about two agents, Caspar and Speedo, who perform the same series of actions with equal success. Caspar, however, takes a deliberative and reasoned decision making approach to each, whereas Speedo proceeds in an impulsive and reckless manner. In the end both are successful, but the argument is that Caspar's approach is more likely to yield success in the long run (i.e. more often) than would Speedo's approach. In essence, the normative justification of formalised Risk Management rests on a claim of improved and more reliable performance achieved through the practice of comprehensive rational planning.

Indeed, at a conceptual level, rational choice implies the assessment and control of risk (and vice versa). That is, decisions involve risks (both upside and downside), which must be evaluated rationally, and those risks involve decisions about controls, which must be made rationally. The Risk Management process is therefore an integral part of making good (rational) decisions, while also involving (rational) decision making in its application. With Enterprise Risk Management, this general rationale is extended to the organisation as a whole: if formalised Risk Management improves (or should improve) the performance and reliability of any particular system or process then the application of formalised Risk Management to all systems and processes in an organisation should improve the performance and reliability of the enterprise at large. In the ideal, when ERM is fully developed within the firm the Risk Management process is said to be proficiently and

consistently performed as a natural part of all processes and roles in the organisation (i.e. “fully integrated”), the process is formalised and explicit in its execution and supported by specific tools and calculative infrastructure, and there is an explicit role for science and evidence in decision-making (MacGillivray et al. 2007a, 2007b; Pollard et al. 2004; Ward 2003; 2005; see the overview of the features of mature ERM in Appendix V). Interpreted in this way, the various ERM frameworks, such as COSO, AS/NZS 4360, and ISO 31000 prescribe a formalised procedural system for making “good” decisions (and for demonstrating ex post facto that those decisions were aligned with the corporate objectives).

The proposals put forward by Watercare’s CRM for developing the company’s risk management capabilities, and his focus on the quality of the company’s “risk data” in particular, reflect an interpretation of his role in these terms. The CRM’s proposals were consistent with the distinguishing features of international models of ERM, and with the calculative ideals expressed in engineering guidance for risk assessment. The CRM proposed to intervene within the organisation to induce organisational actors to follow the precepts of the normative Risk Management process, and to provide procedural and analytical aids to assist them in doing so (i.e. risk assessment frameworks and registers, and methods and tools for risk quantification).

The CRM was especially motivated in this regard to ensure that the organisation’s decision making was supported by objective information about risks (i.e. “good” risk data), where “objective” can be understood in at least two senses. First, in most contexts objectivity means fairness and impartiality, and is considered essential for justice and honest government, and the good governance of organisations (Porter 1995). In this sense objectivity refers to the absence of personal or political bias in the making of judgements, to the achievement of a state of neutrality or disinterestedness (Megill 1994; objectivity in *The Blackwell Dictionary of Sociology*; *The Dictionary of Human Geography*). The CRM expressed an appeal to objectivity in this sense through his aspiration to make decisions defensible and to constrain as far as possible the subjectivity of the risk assessment process. Second, objectivity also refers to that property upon which we may judge the validity of knowledge claims, typically associated with ideas such as reality, truth, and reliability (Megill 1994). In this sense, the term relates to a subject (i.e. a person) who perceives and an object which is perceived. Objectivity is held to be a property of those perceptions referring to the degree to which they actually correspond to a true condition of the world; subjectivity in contrast referring to the possibility of perceiving things not as they really are (objectivity in *The Blackwell Dictionary of Sociology*). The appeal to objectivity in this sense can be seen in the CRM’s pursuit of the engineering ideals of systematic, comprehensive, quantitative risk

analysis, and in his aspiration to achieve a high degree of fidelity between representations of Watercare's risk profile and the company's real risk exposure.

On the basis of the foregoing analysis, the following assumptions may be inferred as underpinning the CRM's vision for developing ERM at Watercare. These were:

- That Risk was something real, specific, and well-defined that could, and indeed should be calculated with objectivity and precision.
- That the dual functions of ERM (i.e. governance support and decision support) could be fulfilled through the production of "good" risk data. The CRM expressed this assumption in the idea that good risk data would provide both insight for decision making and a defensible record of decision making after the fact.
- With regard to the decision support function, that the quality of "risk data" was positively correlated with knowledge about risks and hence with the quality of decision making. On this basis, better quality data would be associated with better decision making, and, according to the general assumptions of ERM, with improved enterprise performance.
- That the task of implementing ERM was essentially one of "design and build". That is, the CRM's approach reflected an interpretation of ERM as essentially a new capability that needed to be integrated into the organisation by literally designing and building the necessary infrastructure (i.e. risk frameworks and registers, and calculative tools), and training organisational actors to perform "Risk Management" (i.e. creating the so-called "risk culture").

In this regard, the CRM's early approach to his role may be characterised as consistent with what Mikes (2010) called the "strategic controller" archetype of the CRO role (see Chapter 1). Chief Risk Officers of the "strategic controller" persuasion approach the role with a focus on integrating risk and performance measurement, and ensuring that risk metrics are deemed reliable and are relied on for decision making purposes (Mikes 2010). The approach is characterised by what Mikes referred to as *quantitative enthusiasm*, a commitment to extensive risk modelling, a belief in such models as robust and relevant tools in decision making, and a primary objective to measure the aggregate risk profile of products and business lines (Mikes 2010). The CRM's approach to his role at Watercare exhibited all of these features.

Although the CRM's approach would later be challenged, it was not unreasonable in light of the CRM's background and the lack of guidance for the implementation of ERM in the

international literature. International frameworks and standards for ERM outline the main objectives and components of an ERM programme, but as conceptual documents they provide little guidance for the implementation of specific elements in specific organisational contexts. Chapter 1 argued that this “method gap” is not restricted only to those broad framework documents, but is a feature of the ERM/CRO knowledge base in general. That is, despite the fact that ERM and the CRO role have been evolving for nearly 20 years, there is as yet very little robust (i.e. theoretically and empirically grounded) guidance for how to achieve what the CRM would later refer to as “the holy grail” of integrated Risk Management (see Dialogue 13). It is left up to Chief Risk Officers to interpret how the ERM concept should be translated into practice.

However, it is easy to interpret broad ERM frameworks, and the corresponding capability maturity models, as implying that “integration” means formal risk assessments should be performed everywhere (i.e. in all business activities), and that the primary goal of those assessments should be the production of what the CRM referred to as good “risk data” (i.e. the pursuit of objectivity). This interpretation is especially intuitive if the CRO in question is of the “strategic controller” type and therefore already predisposed to interpreting Risk Management as primarily an exercise in quantifying probabilities (both the CRM and I initially fell into this category; a disposition which may be attributed to our engineering backgrounds). In other words, in the absence of specific guidance there is little to dissuade CROs from interpreting the task of implementing ERM as one of defining risk policies and frameworks, designing risk registers, and promoting quantitative methods for risk calculation.



# Chapter 4

## Experiences with designing ERM infrastructure

CRM: What I'm trying to do is to provide the parameters that you guys would consider to be your bottom line, the things that are fundamentally important to your business operations...

WWM: Well... I know what my ten greatest risks are and this doesn't change them. Is this being used to try and make them more transparent?

*An exchange between the CRM and another manager over the CRM's new Operations Risk Framework (Dialogue 4, p 270)*

The CRM's vision for developing Watercare's Risk Management capabilities was dominated by proposals for improving the quality of the company's "risk data". Chapter 4 describes and analyses the CRM's experiences as he engaged in certain tasks toward the implementation of that vision: (i) redesigning the company's risk framework and registers, (ii) promoting quantitative risk modelling, and (iii) facilitating strategic risk assessment. A detailed account of the CRM's work is presented in Appendix VI, which constitutes the main empirical data set for the analysis in this chapter. The main points of argument in the chapter concern the descriptions and explanations of what happened in each case, what those experiences reveal about risk as an object of inquiry, and how those insights, in turn, explain the CRM's experiences. The chapter also draws lessons about the design of ERM infrastructure from the CRM's experiences.

Empirically, Chapter 4 sets the scene for Chapter 5 by revealing how the CRM was relatively unsuccessful in his efforts to improve the quality of Watercare's "risk data" through

the design of ERM infrastructure, as he had initially envisaged, but did achieve success by actively facilitating and guiding inquiry. Analytically, the discussion in Chapter 4 lays the groundwork for that in Chapter 5 by (i) framing up the conventional theoretical distinction between Risk and Uncertainty, (ii) explaining the CRM's failures as the result of a conflict between expectation and reality over how detailed "risk data" should be, and (iii) explaining the CRM's success as the product of his alternative facilitative approach.

## **Redefining Watercare's risk assessment framework**

The CRM perceived that many of the problems with the quality of Watercare's "risk data" stemmed from problems with the company's risk assessment framework<sup>6</sup>. The redefinition of that framework was therefore one of the first tasks which the CRM undertook and one to which he devoted considerable attention. The following discussion explicates the key themes from those experiences.

### **The design rationale**

The CRM's primary objective in redefining Watercare's risk framework was to improve the quality of the company's "risk data". As the CRM used the term, "risk data" referred to the information contained in the corporate risk register, representing the outputs of risk assessments performed by staff. The function of that "risk data" was ultimately communicative, i.e. to communicate perceptions of risk between different stakeholder groups in the organisation (most obviously between staff and management, or as the CRM referred to them, between the "coal face" and the "Board room"). From very early on, the CRM had recognised that the risk framework had a normatively important role to play in providing a common language for this communication of risk (see Chapter 3).

At Watercare risk was one of two primary criteria on which capital decisions were based (the other being financial cost). In order for any actor to secure capital funding for a project, he or she was required to represent the need for and benefits of that project in terms of risk. The company's risk framework ostensibly defined the common language and format for that

---

<sup>6</sup> In this chapter the term "framework" refers explicitly and only to the physical document which specified the categories and evaluation criteria by which risk was to be assessed within the company (at Watercare this was a table with five consequence categories and six levels of severity). This is a rather narrow usage of the term "framework" in contrast to its broader meaning which also encompasses policies, rules, and procedures for risk management (see Appendix 1 for a broader description of Watercare's risk management framework).



representation, and was, thus, an obligatory point through which all requests for capital and operational expenditure were required to pass (Latour 1987, 1988). Through this function the risk framework was part of the infrastructure for the governance and internal control of the organisation. That is, the risk framework acted to enforce the maxim that capital funding should only be allocated to projects which contributed to the achievement of the corporate objectives. The CRM expressed this governance role as follows:

...I've realised that one of the main reasons we use risk data is so that people can make an assessment of business need considering the corporate objectives. So, at least in theory, you can look across the entire enterprise and say "this risk data gives me a basis to decide what to do". It's an objective method for determining the priorities based on what the organisation wants to achieve, not what I want to achieve. So, to me, you have to have a consistent and objective method for calculating risk, otherwise you can't compare across the organisation. (Dialogue 8, p 301)

The fact that different elements of the business pursued incommensurate objectives (i.e. water supply vs wastewater treatment and disposal), but competed for the same limited pot of capital funds, emphasised the need for the framework to objectively represent the corporate objectives and performance standards:

...what you're saying is that the corporate objectives will in effect define the relative importance of different things. And that's something I'm very conscious of at Watercare. If I develop the risk framework using the objectives currently under the Statement of Corporate Intent, then I think it would mean that we would stop spending money on wastewater, which would be completely unpalatable to [the wastewater side of the business]. I mean, in essence I'm saying that all risks will be assessed on the significance to the achievement of corporate objectives, but the objective most directly relevant to wastewater is that there will be "no successful prosecutions under the Resource Management Act". The problem is that it says "prosecutions", it doesn't say "compliance with resource consents"... [and] the probability of us getting successfully prosecuted for wastewater overflows is probably very small. In which case, why spend money on it? I realised, holy shit, if we asked [the Wastewater Treatment Plant manager] to do a realistic assessment of the chance that we will get prosecuted, its very small, and his risks will come out Class 2, and he'll never get any budget. (Dialogue 8, p 301)

During his initial review of Watercare's risk management function, the CRM developed the opinion that the company's existing (2003) risk framework was not sufficiently well defined to properly fulfil this boundary infrastructure role. The arbitrary categories and vague criteria within the framework promoted rather than constrained the subjectivity of the risk communication process. To address this, the CRM envisaged that the new risk framework should be, as a far as possible, an objective representation of the company's objectives and

performance standards (this rationale is represented in Figure 4.1, overpage). He conceived of an idealised hierarchical prioritisation, cascading from the company's strategic objectives to the specific performance standards and criteria against which the company was held accountable in each of the various functional contexts of the organisation (see Chapter 3). The CRM referred to this as marrying together the strategic (top-down) and operational (bottom-up) perspectives of the organisation (see Dialogue 7, p 295, & Dialogue 9, p 305), which he felt was essential to the design of a good risk framework.

In this regard, the framework would also serve as an external arbiter of the value judgements inherent to the concept of risk, because a rigorously defined framework would embed prior decisions about the relative priorities of objectives (and hence of risk outcomes), rather than leaving those judgements up to the subjective perceptions of individual staff members. The CRM's approach was therefore to limit the subjectivity of the risk framework in use by seeking clarity and detail in the design of the framework (see Figure 4.1).

### **The art of marrying abstract objectives to concrete practices**

The CRM initially set out to identify, examine, and distill a considerable volume of documentation in order to identify the primary sources of Watercare's corporate objectives and performance standards and criteria. This research included: the Local Government Act and the Company Constitution, Watercare's Statement of Corporate Intent, Annual Report, Asset Management Plan, and Funding Plan, the bulk water and wastewater agreements with the LNOs, the New Zealand Drinking Water Standards and Ministry of Health guidelines on the grading of drinking water supplies, and a review of statutory obligations and compliance penalties in the areas of health and safety, human resources management, resource management (consents), and financial reporting.

In aggregate, those documents specified a broad range of objectives at each level of the administrative hierarchy, a plethora of performance standards (both internal and external) to which the organisation was held accountable, and hundreds, even thousands of parameters for measuring every facet of Watercare's performance. Furthermore, those standards, together with the processes and instruments for measuring and reporting performance, constituted a complex metrological system which was not the product of a co-ordinated, top-down design, but rather reflected the historical and ad hoc development of the company and its technologies and practices, as well as the evolution of the broader industry and regulatory environment in which it operated. The risk assessment framework envisaged by Watercare's CRM sought to represent that complex metrological system within a highly simplified matrix.

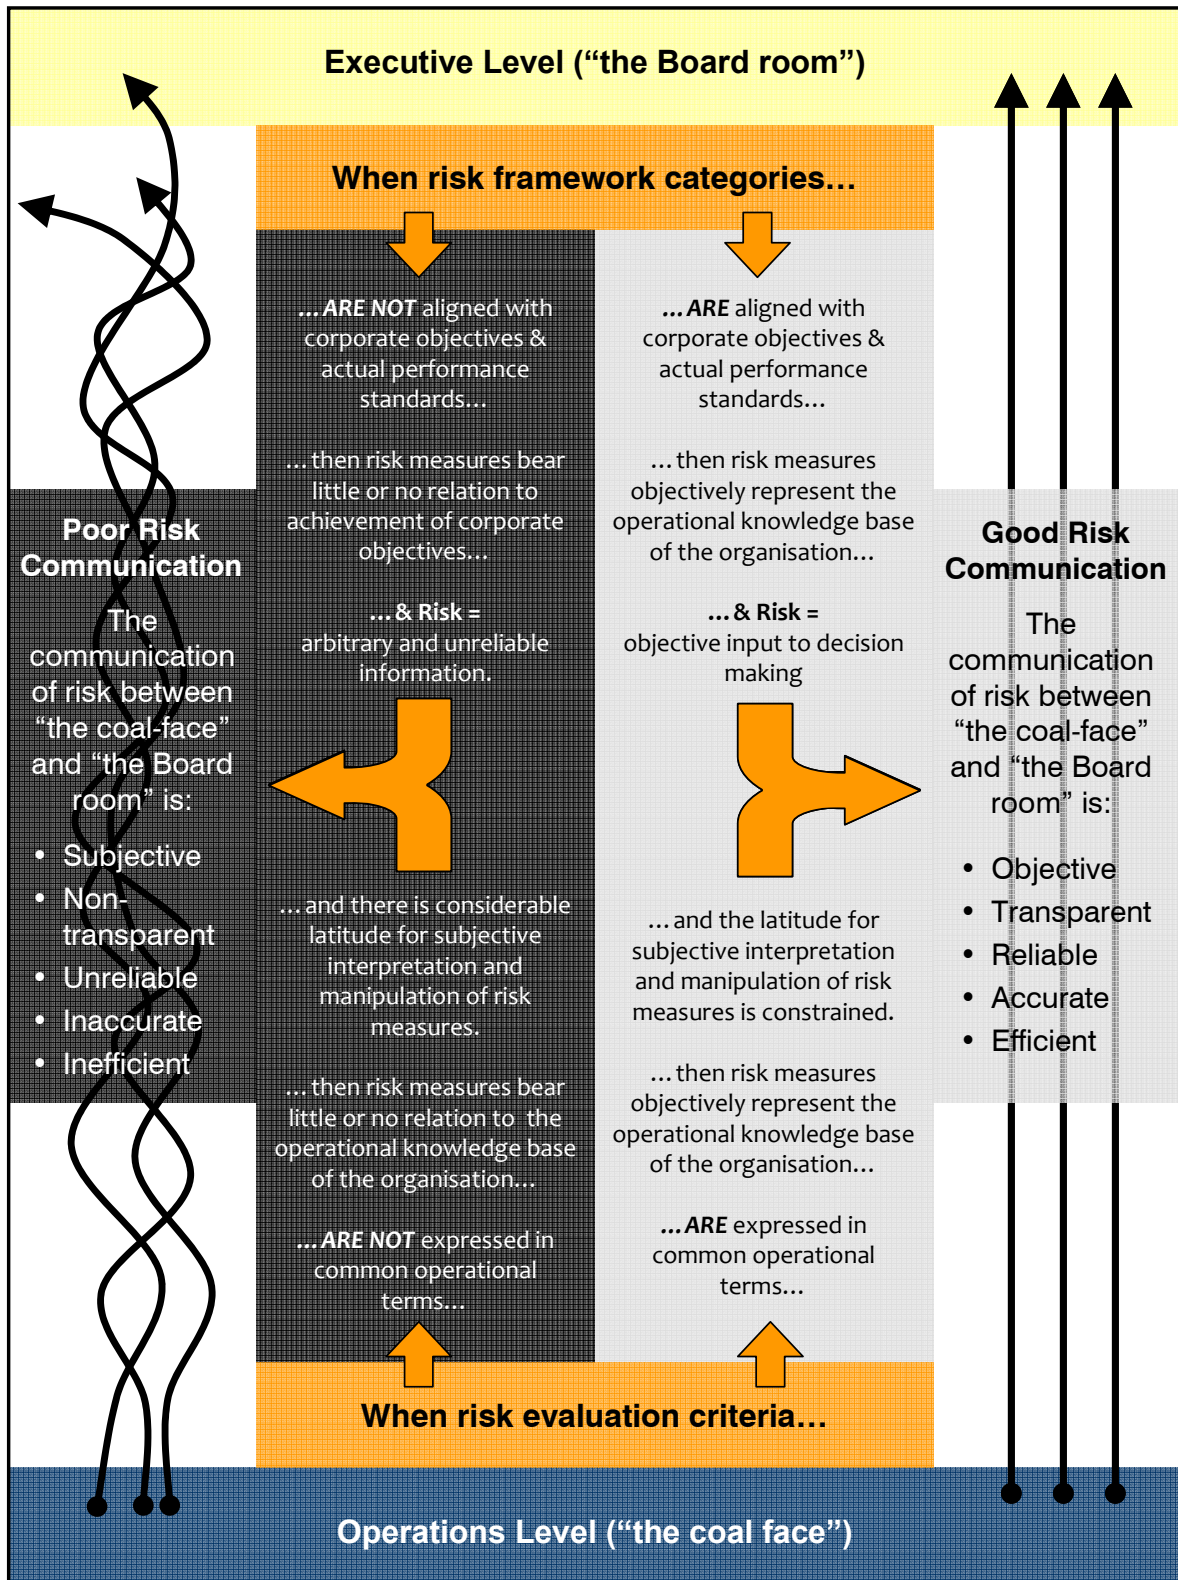


Figure 4.1. Illustration of the CRM's rationale for how the structure and content of risk frameworks can affect the representation and communication of risk between communities of practice in the organisation: the proper alignment of risk framework categories to corporate objectives and performance standards, and the proper specification of risk evaluation criteria in common operational terms, are assumed to improve the objectivity of risk communication.

Since it is the nature of the task of simplification that something must be left out, the design of that framework necessarily involved decisions about which objectives to represent and their priorities, which performance parameters to use, and how to organise both objectives and parameters into a coherent structure which objectively represented “risk” consequences for the organisation.

The CRM discovered, however, that the answer to the question of “which objectives are most important?” was not straight-forward, even for someone with a good general knowledge of the organisation. Indeed, there were different answers depending on who answered the question (see, for instance, the CRM’s engagement with Operations line managers in Dialogue 4, and the analysis of that dialogue in Appendix VI, p 472-474). This is not to say that organisational actors are free to arbitrarily prioritise whatever they want, but rather that in any context there may be multiple legitimate ways of prioritising “what matters most” with respect to objectives and performance. Indeed, the fact that the CRM’s revised risk framework would be just one of at least four different enterprise-level systems for the evaluation and representation of Watercare’s performance developed within the company since 1999 pointed up the multiplicity of ways in which the organisation could be interpreted and represented (see Appendix VI, p 480). The CRM thus acknowledged that the task of designing the risk framework was one of trying to “boil down” the “raft of possible ways you could measure organisational performance” to “a handful of those that matter” (Dialogue 13, p 325).

The CRM could not resolve the question of what mattered most on his own; or at least not solely by a detached deconstruction of the statements and criteria contained in corporate and regulatory documents. Rather, the task of sorting through and prioritising the complex system of objectives and performance parameters applicable to the various functional contexts of the organisation required considerable operational knowledge of those contexts. This was revealed in all three of his engagements with the stakeholder groups in the Watercare organisation (these engagements are described in detail in Appendix VI):

- Through his conversations with various managers from the Asset Management group the CRM discovered that the Planning and Project Management units generally had relatively little impact on the immediate (day-to-day) performance of the Watercare enterprise, but were fundamentally responsible for the company’s long-term performance against its primary legislative objective. In this regard, while project-level risks were, individually, often of relatively minor significance at the enterprise level, the systemic risk posed by consistent under-performance of those business units was

considerable. It was this knowledge which emphasised to the CRM the importance of tying project-level risk assessment to the performance of individual project managers.

- The discussion in the August 2007 meeting with the Operations line managers revealed the complex range of inter-related standards to which the company was required to perform, particularly with respect to the delivery and treatment of reticulated water. For instance, for water supply there were performance standards relating to factors of immediate health significance (P1 Determinands), factors of long-term health significance (P2 Determinands), aesthetic factors (taste and odour), the quality of management systems and personnel (MoH Grading), water flow and pressure, and drought security. Evaluating the relative significance of breaches of these standards involved consideration of the importance of the standard (e.g. P1 vs P2 Determinands), the magnitude and duration of the breach, the number of people affected, control actions that the company might take (e.g. imposing water restrictions, issuing Boil Water Notices), and subsequent outcomes (e.g. people getting sick, media attention, prosecution of company personnel, restructuring of the company). The CRM was only able to sort through this plethora of parameters and structure a coherent framework by drawing on the detailed operational knowledge of the actors at the “coal face”.
- Through his engagement with Watercare’s General Managers the CRM discovered that, while the Statement of Corporate Intent was officially the public statement of Watercare’s corporate objectives, in practice the SCI was strategically irrelevant. This was not due to wilful disregard of the document by Watercare’s General Managers, but was rather a product of their in depth understanding of Watercare’s institutional, regulatory, and physical operating environments. In that context, the SCI was important as a statement of Watercare’s corporate social responsibility, but had little impact on the strategic development of the physical water and wastewater infrastructures. The primary forces in that regard were Growth (changes in demand), Levels of Service (changes to regulatory and contractual requirements), and Renewal (replacement or rehabilitation of aging assets). It was this understanding which led the CRM to refocus his attention on the statements of Watercare’s principal objectives in the Local Government Act and the Company Constitution as the appropriate starting point for deriving his risk framework.

The common feature in each of the above engagements was that the CRM elicited important understandings about “what mattered” in terms of performance in each of the

functional contexts:

...from a personal point of view, I've always heard that historically the guys who were really effective at running companies were the guys who could walk down to the workers on the shop floor and ask 'what do you do here, what are your needs, what stops you doing what you need to do?' Which is what I've tried to do, and I guess my job then is to go away and try to link that with the corporate objectives. So I think the value that I add is in understanding what they need at the coalface in relation to the corporate objectives, and being able to put it all together in an efficient and consistent structure (Dialogue 3, p 257)

Prior to his engagement with the various stakeholder groups, those understandings were not readily apparent to the CRM as an outsider, even though he already possessed considerable knowledge of Watercare's corporate objectives and performance standards. Individual actors' perceptions of significance were shaped by a range of factors specific to their practice contexts, including detailed operational understandings of the technical and organisational systems, the nature of the physical outcomes arising from failure of those systems, expectations about public and political perceptions of performance failures, and expectations about personal accountabilities. These were quite specific understandings about how things worked in each context, and about what was important and what was not, acquired by the various actors through long practical experience, both in their professional fields in general, and in their specific roles within the company. It was only with these understandings that the CRM was subsequently able to sort out which objectives and performance parameters should be represented in the framework and how they should be organised.

Through this process the CRM also established connections between the objectives and parameters in the risk framework and existing capabilities for data capture, analysis, and reporting within the company. The CRM believed that such connections justified the use of certain parameters over others because they could be objectively quantified through existing business processes. Dialogue 9, for instance, contains a description, in the CRM's own words, of how he could have interpreted ("sliced") the corporate objectives in different ways, and of how he focussed on a particular solution because he wanted to link the framework to Watercare's existing analytical (i.e. modelling) capabilities:

Always in the back of my mind was this thought that we need something that we can model, we need to have a measure of consequence that we can model, and since we already hydraulically model the water networks I wanted that to be a measure of consequence... The way that I'm slicing and dicing the objectives, or rather the way that I'm analysing the organisation will enable us to use computational models to do some of the work. (Dialogue 9, p 304)

For the CRM, the possibility of using Watercare's existing models and data to quantify parameters in the risk framework was an important measure of the degree to which he had succeeded in the task of marrying together the strategic focus with the operational focus (just as it was also an indicator of the degree to which the existing framework failed in this task):

I think it comes back... to the fact that I've been trying to marry the operational focus with the strategic focus. So I started with the objectives and tried to break them down, but I always thought that at some point I'd love to find that they marry well with what the guys on the shop floor think about. So I've talked to them and worked out what's important to them, the point really being that if we're doing computational modelling then its probably computational modelling of what they think about, isn't it? (Dialogue 9, p 305)

In effect, the links to various practical tools and capabilities, and the contextual understandings revealed through the CRM's investigation, constituted objective support for why certain objectives and performance parameters should be represented in the framework instead of others. So, for instance, the existence of certain computational models within the company was indicative that the parameters quantified by those models were sufficiently important to warrant the development of special infrastructure for their calculation. Thus, the company's existing calculative capabilities, and the managers' specific knowledge of their functional contexts, subsequently supported and justified the CRM's framework in the face of counter claims that alternative (i.e. extant) definitions of the risk framework might be better.

For the CRM, the process of (re)defining Watercare's risk assessment framework was thus revealed to be less a detached exercise in deconstructive logic and more an artful process of blending, into an explicit text, certain understandings about what was important and how things were done in particular contexts. The understandings were not the CRM's but were, rather, those of members of the various communities of practice (general managers, planners, project managers, operations staff) which would ultimately have to use the risk framework. In Dialogue 13 the CRM described this process as follows:

...coming up with that framework is really part art, part science. You need to be able to look at a business and understand it, and at the same time you need to be able to look at your objectives and deconstruct them, and then try to marry the two together. You have some paper-based objectives, and then you have real business practices and what you're trying to do is to bring these two together to achieve an efficient translation of information from the bottom to the top. (Dialogue 13, p 327)

This translation between strategic objectives and operational parameters is reflected in the revised Operations risk framework which the CRM eventually proposed (described in detail in Appendix VI, page 489).

## Dilemmas

The CRM's pursuit of an objective risk framework gave rise to two dilemmas. The first was that of the flexibility of a detailed risk framework relative to changes in the corporate objectives, performance standards, and parameters. The CRM was proposing a number of inter-related frameworks for assessing risk in different functional contexts of the organisation (see Appendix VI, p 466). The idea was that the more specific the framework was to a particular context, the more effectively it would constrain the subjectivity of risk assessment and communication between staff and managers in that context, and the more objective would be the resulting "risk data". However, as the CRM's experience showed, the more detailed and objective the frameworks the more difficult and time consuming they would be to produce. Indeed, it took the CRM the better part of 12 months to propose a redefined Operations risk framework<sup>7</sup>. It further occurred to the CRM that, having gone to the trouble of defining those context-specific frameworks and the relationships between them, if the overarching corporate objectives or the specific objectives of individual business units subsequently changed, the frameworks would have to be redefined (see Dialogue 8, p 300). The same issue arises with the performance standards and parameters referred to by the risk frameworks. If those standards and parameters change, then, again, the frameworks would have to be adjusted to reflect those changes. This problem of changing objectives and standards was not of major concern in Watercare's context, where rates of change in the institutional and regulatory milieu were relatively slow, but it was reasonably foreseeable that it could be a significant problem in more dynamic environments. This observation gave rise to a dilemma: on the one hand, according to the CRM's rationale, a detailed risk framework was desirable for objectivity, but, on the other, the more detailed the framework, the less flexible it would be. The CRM also commented that too much detail might result in the framework always being out-of-date, leaving the framework as merely a backwards looking statement of what was (Dialogue 8, p 300)<sup>8</sup>.

The second dilemma was that a detailed risk framework might prove to be too complex. The CRM had already commented in Dialogue 1 that he felt constrained by what he

---

<sup>7</sup> In this regard, the design of risk frameworks for public sector organisations (local authorities, government departments, etc.) may be more of a challenge than for private sector firms due to the strong mission foci of the former (i.e. in the private sector overriding emphasis is usually given to financial performance, whereas it is just one amongst a number of core mission objectives in the public sector).

<sup>8</sup> This also highlights that, as an instrument of rationality, a risk framework can only be defined in the context of given objectives. The question of what the enterprise should aspire to is, by definition, beyond the scope of risk management. This is, rather, a normative question for consideration by the firm's stakeholders.



perceived to be a cultural and practical preference within Watercare for “simple” risk management (see below). The problem was that the rigorous definition of an integrated risk register reflecting the functional and informational hierarchy of the organisation would result in a large number of risks, i.e. as higher level risks are refined into more detailed sub-risks. The CRM commented as follows:

...one of my overall objectives is to simplify the register, and introducing an additional 300 risks would just go completely against that philosophy. Now I want to get better insight into the data but doing that by a mathematical method would deliver a huge data set that would be too cumbersome to manage.. Also, there’s a perception inside the organisation that the number of risks in the register is proportional to the efficiency of the system, and when I’ve talked to people one of the primary factors that influence’s their level of buy-in is the amount of extra work that’s imposed upon them by the risk management function. So given that we’ve got eight hundred risks at the moment and the system is so cumbersome, and people are talking about having just a hundred risks, well, I’m not going to be able to sell this. (Dialogue 1, p 242)

The same issue was at play with respect to the development of the risk framework. That is, one effect of a detailed framework would be to force staff to define risks in greater and greater detail. To achieve this, staff would have to both possess and express relatively detailed knowledge of the risks they were assessing, preferably in quantitative terms. It was easily perceived by staff that such a detailed risk framework would impose an additional expenditure of time and effort to both understand the framework and perform risk assessments with it:

...I mean Watercare has an AMP process which prioritises the company’s projects. At the moment that prioritisation is all pretty much intuitive, but effectively that process is already doing what I’m trying to capture here. I’m just trying to put a more formal assessment process on it, to say that this objective is actually more important to the organisation than this one. It’s an awkward question for the organisation to deal with, and even do you need to, I think, is the bigger question. I mean, that’s what’s rolling around in the back of their eyes. There’s a perception on their part that I might be trying to put mathematics onto something that they do already. So to be honest it’s a fair question. Do you need to do this when there is already a process in place to do it? When does it become too much? (Dialogue 9, p 304)

The dilemma for the CRM was that while a detailed risk framework was desirable for objectivity, it might also be rejected by staff and management as too cumbersome to be used as an effective management tool.

## Difficulties using the new Project risk register

After redefining Watercare’s risk assessment frameworks the CRM undertook to redesign the company’s risk registers. This was necessary to ensure that the registers were consistent with the new frameworks. Appendix VI (pp 483-488) describes the CRM’s approach to redesigning the project risk register, and the surprising difficulties that were encountered when it was put to use in a project risk identification workshop (the Hunua No. 4 workshop<sup>9</sup>). The CRM’s objective with the project risk register was to improve the specification of risk information, in particular to facilitate clarity and consistency when describing and linking together the causes and consequences of identified risks. The difficulties experienced with the new register in use raised a perplexing question: why had such a carefully designed tool proven so difficult to use in the very context for which it was intended? The following discussion reflects on this and subsequently raises questions about one of the fundamental assumptions underpinning the CRM’s approach – that the quality of the company’s “risk data” was a function of the level of detail in the risk assessment framework.

### Specifying risk requires a certain degree of knowledge

The new project risk register required the user to specify the particular objectives for the project in question (functional output criticality, budget, schedule; see Appendix VI, p 483). This information was subsequently used within the register to calculate the significance of project risks within the project and corporate profiles. This was in contrast to the existing register which required those impacts to be evaluated against an essentially arbitrary severity scale originally intended to evaluate significance only at the enterprise level.

The participants in the workshop sought, through the evolution of the discussion, to develop understanding around the following questions (even if these questions were not necessarily made explicit):

- What project component(s) are we talking about?
- What are the specific performance objectives/requirements for this component in light of the overall objectives for the project? (i.e. how does each component

---

<sup>9</sup> The Hunua No. 4 project was to plan and construct a large diameter treated water main approximately 20km in length (the Hunua No. 4 main). The new main was required to meet long term demand growth and also for operability reasons so that existing mains could be taken out of service for maintenance. The estimated capital cost of the new main was NZ\$195 million.

contribute to the achievement of the overarching project objectives?)

- What do we know about this component? (e.g. how will it be constructed? what has to be done to produce it? etc.)
- Given what we know, what might prevent those objectives being achieved, and/or what opportunities are available to realise efficiencies in achievement of those objectives?

The difficulty revealed in the Hunua No. 4 workshop was that definitive answers to the above questions were not forthcoming for several of the project components under discussion. Rather, the discussion focussed on the nature of the uncertainties, and what needed to be done to resolve them; understandings which only became concrete (and subsequently amenable to written description) in terms of their scope and implications as the discussion evolved. Even then, however, what the specific risks were, if any, remained unclear.

The experience in the HN4 workshop reinforced the fact that no matter how carefully the overarching project objectives may be specified they simply declare the performance requirements for the project as a whole. They do not say anything about the specific performance requirements of the individual project components, of which there are usually many, and which combine in varied and complex ways to achieve the overarching objectives. Generally speaking, individual components will be located in different places (sometimes in close proximity, sometimes distributed over a large geographical area), and the specific performance requirements for each component will be very different (e.g. the performance requirements for a bridge are very different to those of the pipe lining or a pressure valve). The location and function of each project component defines its specific performance context. For instance, the HN4 pipeline was over 20km in length and the performance requirements of a sub-component called “the Pukakai River Crossing” were quite different to those of another sub-component called “the Mangere Main Cut-in”, which, again, were quite different to the performance requirements of the “pipe lining” (to use examples from the workshop).

This suggested that in order to define a specific statement of risk for any individual project component it was necessary to define context twice, once broadly, and once specifically. The overarching objectives for the project (i.e. schedule, budget, and deliverable) defined the global criteria for assessing the consequences of risk events, and were defined once for the whole project. But, excluding general “catch-all” risks, specific risk statements are made about, and in terms of, a particular element or component of the project. Thus, the

specific objectives for each project component, and the specific systems and processes that would generate those components, had to be defined if the component-specific risks were to be identified and described in clear and unambiguous terms. That is, the person describing and assessing the risk had to know these things, or have the relevant information at hand, in order to construct the statement of risk.

## **Knowledge of risk has to be produced**

My experience attempting to capture the outputs of the HN4 workshop using the new Project risk register highlighted an important point: that, like all knowledge, knowledge of risks is an evolving and emergent process (Buenaño 1999; Cook and Brown 1999; Nonaka 1994; Tsoukas and Vladimirou 2001). In particular, to document risks in the form required by the new risk register (i.e. with clear differentiation between causes and consequences, quantified measures of consequence, and specific controls clearly addressing either likelihood or consequences) required important distinctions to be made, distinctions which only become possible toward the end of what might be a fairly lengthy and involved process of inquiry.

This reflects the normative conception of the risk assessment process as a step-by-step process of inquiry which seeks to answer several questions (Kaplan and Garrick 1981; see Appendix V): (1) What are the outcomes (consequences) of interest? (2) In what ways might those outcomes be realised? And, since the future is uncertain, (3) how likely is it that those outcomes will actually occur within the given time frame? The point is that it only becomes possible to describe risks when these questions have been answered. In the absence of “sufficiently complete” knowledge of the system objectives, of the system components and their functions and relationships, and of relevant external variables, it becomes increasingly difficult to specify risk. Indeed, there is always the possibility that the inquiry may not succeed in producing “sufficiently complete” knowledge, depending on the time and resources available, the complexity of the system, and the intellectual capacity and experience of those performing the inquiry (Buenaño 1999). In this sense, making a clear and specific statement of risk is a highly evolved capability that is acquired through processes of inquiry (in this case the process of planning and design). Risk identification workshops represent only one element of that process.

The difficulty in clearly defining project risks in the HN4 workshop was, therefore, not due to any lack of competence on behalf of the participants, or even the lack of structure in the discussion, but can be more productively explained by the timing of the workshop within the project programme. That is, due to the state of the project design at that time,

considerable uncertainties remained with respect to various project components and it was this lack of certainty which prevented us from formulating statements of risks in specific and quantitative terms. This suggests that workshops held at different stages of the project life cycle will be more or less successful in defining project risks to a high level of detail due to the degree of knowledge available. In the early stages the specific project objectives (schedule, budget, functional specifications) may not be well defined and the available knowledge of the project system (e.g. modelling data, stakeholder requirements, design of project components, etc.) may be quite uncertain. Under these circumstances the focus of attention will necessarily be on identifying and resolving those knowledge uncertainties<sup>10</sup>. But any expectation that workshop participants could describe such uncertainties in detailed risk terms would be difficult to fulfil, and potentially premature, since adequate knowledge of the issues would not be available.

### **The risk register as a technology of bureaucracy**

Upon reflection, the difficulty that I experienced in using the new risk register was not because the CRM had failed in his goal of designing a register to accurately capture “risk data” in the project management context, but rather that the register had been employed in a situation where “risk data” did not yet exist. It was simply ill-suited for recording the uncertainties which predominated in the Hunua No. 4 workshop.

This apparent mis-match highlights an important limitation with respect to the role of risk register. By its very definition a risk register is simply a tool for inscribing the final output of the risk assessment process in a certain format – i.e. statements of risk. As such, the register ostensibly provides an accurate record of information used in corporate decision making, which serves at least three important functions within the organisation. First, the register serves what the CRM referred to as the “defensibility” function of Enterprise Risk Management, providing at least in principle a transparent record for the purposes of *ex post facto* justification (e.g. under audit). Second, it serves as a database of controls which can be monitored by the internal control function of the organisation. And, third, as a template for

---

<sup>10</sup> After observing several workshops (not all of which were risk workshops) it became apparent that they were less opportunities for resolving uncertainties than they were forums for raising questions about where knowledge may be incomplete. Five general strategies for handling uncertainty were subsequently observed: actors (a) use sensitivity analysis to clarify whether the uncertainty is significant and therefore worth worrying about, or if it can be ignored; or (b) engage in knowledge seeking activity, i.e. further investigations and analysis; or (c) factor out the uncertainty by reframing the problem in such a way that it is not necessary to deal with it; or (d) ignore the uncertainty by making relevant assumptions; or (e) employ all of the above.

recording information in a clear, concise, and commonly understood format without ambiguity, the risk register also supports the risk communication process. That is, the clearer and more precise the statements of risk in the register, the less latitude there is for misinterpretation of those statements by other parties. In all three cases the register functions as an important part of the infrastructure for the good governance and internal control of the organisation.

However, since the ability to make clear and precise statements of risk is only acquired toward the end of the risk assessment process, this would seem to imply (i) that the productive use of the register is necessarily limited to that point, and (ii) the corollary, that the register contributes limited, if any, support to the rest of the risk assessment process. That is, the risk register does not help actors to produce knowledge about risks, it simply records that information once it has been produced.

One could argue that this is not entirely the case since the register does provide support to the risk assessment process as a structure for thinking through the information at hand. Certainly this is true in as much as the categories in the register imply certain questions about objectives, causes, consequences, controls, and responsibilities, but, then, would it not be more helpful for those questions to be made explicit rather than left implicit? Furthermore, as became apparent in the Hunua No. 4 workshop, knowledge about risks does not emerge from the discursive process neatly arranged into structured statements linking causes to consequences. Rather, the emergent information is more likely to be messy and ill-structured, a function of the degree to which the discussion itself is focussed, and the degree of uncertainty which prevails. In practice it can be difficult to rearrange and constrain the emerging information to the pre-defined format of the register on the fly. Indeed, in contexts where uncertainty predominates, where there is insufficient information to answer the questions required by the risk register, attempting to fit that information to the register template can be confusing and frustrating.

The risk register clearly defines the end goal of the risk assessment process and therefore arguably provides a degree of clarity with respect to structuring knowledge and information for communicative purposes. But the further one moves back up the process of inquiry, the greater uncertainty there will be about objectives and system components, the less it will be possible to “structure” that information in strict risk terms, and, therefore, the more limited must be the effect of the register on actors’ subsequent cognitive endeavours. Put simply, the greater the uncertainty, the more difficult it becomes to use the register, and the less actors will find it a helpful tool.

The contention of the above argument is two-fold. First, that as a rigidly defined

inscription device, the productive use of the risk register is necessarily limited to situations where the available knowledge and information about risks is “sufficiently complete” that it can be formatted to the template of the register; that is, to the end of the risk assessment process. Second, that earlier in the risk assessment process, where uncertainty predominates, the effect of the register must be increasingly limited in terms of supporting the investigative and cognitive endeavours involved in the process of inquiry. Thus, in as much as the risk register has an important role to play in how “risk data” is subsequently used, it is predominantly a tool of bureaucracy designed for record keeping and internal control of the organisation, and makes little contribution to the knowledge production process.

This is an important distinction because the production of knowledge about risks (and therefore the risk assessment process) lies at the heart of the risk management process. That is, presumably, one can not manage risks, nor monitor the effectiveness of controls, unless one has already identified and defined those risks and controls. If the ability of organisational actors to do so is not a function of the level of definition in the risk register, then this also calls into question the CRM’s assumption about how the level of detail in the risk framework would influence the quality of Watercare’s “risk data”. That is, since the risk register is simply an inscription template which reflects the structure and content of the risk assessment framework, the above arguments also apply to the risk framework.

The rationale which underpinned both the CRM’s motivation to redefine Watercare’s existing risk framework, and his subsequent justification of the revised framework, was that a more detailed framework, explicitly defined in terms of the company’s actual performance standards, would constrain the subjectivity inherent to the risk assessment process. On the face of it, this rationale seems reasonable. A more detailed risk framework could be expected improve the quality of “risk data” because the resulting statements of risk would be more accurately specified in concrete operational terms. The critical limitation, however, is that by requiring the user to make finer distinctions between objectives, causes, consequences, and controls, the more detailed risk framework raises the bar in terms of the requisite level of information and knowledge that must be on hand in order to construct such accurate statements of risk. In other words, to define risks in the terms of the more detailed framework, users would have to both possess and express relatively detailed knowledge of the risks they were assessing, preferably in quantitative terms. The corollary of this is that the above rationale can only hold where such detailed knowledge is made available, since where uncertainty prevails the user will find it difficult to construct statements of risk in terms of the framework and the apparent objectivity of any statement so constructed will mask the underlying uncertainty (i.e. the resulting “risk data” will only appear more objective). Thus, a

more detailed risk framework implies the need for more detailed risk assessment processes to support the production of “risk data”.

## Limits to modelling the Enterprise Risk Profile

The other core element of the CRM’s vision was the development of quantitative techniques for calculating likelihood and consequence, i.e. risk modelling. This was a response to the CRM’s perception that existing formats used to represent the company’s enterprise risk profile were problematic (see Chapter 3). The CRM envisaged that different types of models could potentially be utilised to quantify the likelihoods and consequences for different types of risks, which would then be combined via Monte Carlo-based simulation to quantify cumulative risk profiles both within and across Watercare’s risk silos. Watercare’s existing reliability modelling capabilities served as both an inspiration and ultimate goal in this regard (see the expression of the CRM’s vision in Dialogue 2).

In system reliability modelling the objective of the calculation is to estimate the reliability of a given component or system in the form of a probability distribution (see the description of reliability modelling at the end of Appendix V). The data requirements for this are significant due to the need to minimise uncertainty. Uncertainty refers to the discrepancy between the true probability of the event and the actor’s estimate of what that probability actually is (Runde 1998; Wright 2002). In reliability modelling, uncertainty arises when the parameters which define the probability distributions for each individual component are unknown, or have to be estimated from a very small sample of empirical data (Dai et al. 2007; Yin, Smith, and Trivedi 2001). In the context of reliability modelling the greater the absence of data the more difficult it is to infer an accurate probability density function for the component in question, and the less certain the user can be that the resulting model will produce reliable results. The problem is exacerbated because aggregate system-level models are built using component-level models, which amplifies the effects of errors and uncertainties in those source models:

...system reliability computed from the function of the uncertain parameters is also uncertain. For large complex systems with many components, the uncertainty of each individual parameter amplifies the uncertainty of the system reliability. Ignoring the parameter uncertainty can result in grossly underestimating the uncertainty in the total system reliability, which, in turn, leads to an overly optimistic expectation of the system reliability and an underestimation of the risk involved when using the system reliability measure for decision making. (Dai et al. 2007, p 783)



The corollary is that the source data for system reliability models has to be at least “good enough” to support the subsequent calculations and simulations (i.e. the confidence intervals for the individual parameter inputs must be very narrow). In effect, this means that in order to have confidence in the system reliability models as objective representations of reality, the scope of those models must be limited to what is known with reasonable certainty (Wright 2002). But a review of Watercare’s asset condition information revealed that it was neither totally comprehensive nor of a consistent quality (see Appendix VI, p 488). It became clear that only a small percentage of the risks identified by the company would ever be specified with sufficient detail that they could be included in probabilistic models. While there were good reasons for this (see Appendix VI), the variability in the quality of “risk data” across the enterprise, ranging from reasonably complete, consistent, and reliable data for some risks, through to virtually non-existent, incomplete, or highly subjective data for others, called into question the CRM’s vision of reliably modelling the Enterprise Risk Profile. That is, the use of probabilistic risk models would have to remain confined to those areas where the company had a comprehensive and reliable data set (i.e. for the water and wastewater treatment plants). Such models would not be reliable in other settings, including at the enterprise level.

## Facilitating the identification of strategic risks

By the latter half of 2007 Watercare’s general managers were becoming increasingly cognisant of the need to update the company’s strategic plan to reflect a range of issues which they perceived would influence the way the company did business in the future. This came through to a certain extent in Dialogue 7. At least two of the general managers were actively working to frame up the plan and saw the CRM’s interest in establishing the ERM programme as an opportunity for the two processes to “dovetail”, as one manager put it. For the Corporate Risk Manager, this presented a significant opportunity to demonstrate the value of the risk assessment process and so, in October 2007, the CRM formally proposed a programme of work to establish a framework for identifying and evaluating strategic risks and to apply that framework for the first time (GM brief on strategic risks, Oct 2007). The CRM’s objective was to collate available knowledge of Watercare’s strategic risks and organise that information within a coherent framework (which would be the new Strategic Risk Framework). The framework would serve as a key input to the strategic planning process due to take place in early 2008.

I participated in this task, helping with data collation and documentation, as well as with

structuring the overarching strategic framework. The CRM's experiences with developing the framework are described in detail in Appendix VI (pp 490-500). The following discussion characterises and analyses the nature of the CRM's performance.

### **Sense-making in a context of uncertainty, ambiguity, and vagueness**

The process evolved as an iterative back-and-forth between (a) extracting information from Watercare's managers, and (b) collating that information into a comprehensible format, with each step informing the subsequent work. A generic framework for categorising information (the PESTEL template) guided the initial collection of information. That data, although too vague to allow risks to be defined, was further rationalised, organised, and extended by the CRM into a more detailed framework explicitly linking external uncertainties to the Watercare system. That work required the CRM to construct a conceptual framework within which it was possible to make sense of the managers' responses.

This sensemaking exercise was far from simple, requiring the CRM and myself to engage in considerable reflection and dialogue over what precisely the notion of a "strategic risk" encompassed in Watercare's business context, and how the 330 or so "issues" that Watercare's managers had identified could be rationalised into a meaningful and manageable set of "strategic risks". The framework and the fleshed-out descriptions of those "risks" subsequently provided a basis for further and more specific data collection.

This took the form of a two-way feedback process. The CRM first fed back the restructured framework and detailed risk descriptions to Watercare's managers in the form of a report which explained how the framework had been developed and which listed the defined risks in the form of a series of templates. The templates summarised the risk assessment results in terms of a number of categories (Clement 2008e, p 21):

- The root source of the risk (underlying causality)
- The nature of Watercare's exposure to the risk (i.e. the extent to which it was avoidable)
- The nature of the expected impacts on Watercare's business
- The degree of understanding about the risk (both causes and consequences)
- Current controls addressing the risk
- Future controls that would need to be developed to address the risk.

After allowing some time for that report to be digested, the CRM engaged the managers in an interactive feedback session in which they were encouraged to share their reflections: Did the framework make sense? What was right about it? What was wrong about it? Were the descriptions of the individual risks appropriate? What was unnecessary? What was missing? What other details could be added? The information collected from that session subsequently informed the further refinement of both the strategic risk framework and the descriptions of the individual strategic risks.

This process was characterised by significant uncertainty, ambiguity, and vagueness from the outset. Any expectation that the CRM may have had that Watercare's managers would be able to precisely and quantitatively describe their perceptions of future trends or conditions "that could have a major impact on how WSL does business in the medium to long term" was quickly dispelled by their initial responses to the PESTEL template. It became clear from the scale and interconnectedness of the issues identified by Watercare's managers that any assessment of potential effects would be at best a "guesstimate", at least to begin with.

In addition, when the CRM initiated the work on the strategic risk framework he was himself uncertain about what the notion of "strategic risk" meant in Watercare's context. His initial concept, as revealed in the question he posed to Watercare's managers, was only loosely defined around notions of "longer term" and "major impact" on "Watercare's business", while the PESTEL template was simply a convenient, if arbitrary, way of categorising the answers to that question. It was not until a month later that the CRM's assumption about the strategic importance of the Statement of Corporate Intent was called into question during a meeting with Watercare's General Managers (see Dialogue 7), and we did not fully make explicit a working definition of "strategic risk" until late January 2008. The former oriented the CRM's attention to the statements of the company's core objectives in the Local Government Act, consideration of which revealed that the company really had just one persistent strategic objective: to maintain a balance between the long-term cost of the company's water and wastewater services and the long-term risks to the delivery of those services (see the explanation of the description of Watercare's mission and objectives in Appendix I).

It was only after some considerable discussion of that objective during the latter workshop that we were able to resolve the components of a formal definition of "strategic risk" applicable in Watercare's context (see Table 4.1). Specifically, the CRM defined a strategic risk as any event or scenario which threatened the ability of Watercare to fulfil its statutory purpose under the Local Government Act, *and* which required the company to change the way it did business in order to avoid, accommodate, or cope with the potential

adverse implications of the identified scenario (Clement 2008e). This compound definition provided the necessary criteria by which to distinguish strategic risks from operational and project risks in Watercare’s context, and to evaluate the magnitude of those risks. Strategic risks were those that required the company to implement a strategic response, being one which resulted in the company developing a new capability. Operational risks were those for which the company already possessed the necessary competencies for control.

**Table 4.1.** The CRM’s definitions of key risk categories  
(source: Clement 2008e, Fig. 2, p 7)

Risk Category	Definition
Strategic Risks	Risks associated with successfully leading / guiding the evolution of business competencies.
Project Risks	Risks associated with successfully delivering new business competencies.
Operations Risks	Risks associated with successful operating business competencies.

The process described above was, for the CRM, one of evolving the managers’ initially vague perceptions into more concrete and reliable information, while simultaneously resolving his own uncertain concept of “strategic risk” into a coherent framework. This was an important difference with the earlier work on the operations and project management risk frameworks where the CRM was concerned only with specifying the categories and parameters by which risk should be defined and quantified. In both of those cases the resulting risk assessment frameworks were essentially abstract reference documents which listed, in organised form, the various parameters that could be used to describe and quantify risk in particular contexts. With the strategic framework, however, the CRM undertook to both define the overarching assessment framework, *and* to facilitate the identification, description, and evaluation of risks at the same time. The two tasks were fundamentally intertwined.

Through the above process, the CRM was able to resolve some of the initial ambiguity and vagueness and to formulate meaningful descriptions of the identified strategic risks. By “meaningful” I refer to the fact that the CRM’s Strategic Risk Report (Clement 2008e) was able to make explicit *specific* causal links between potential changes in *specific* environmental and organisational factors, the potential subsequent effects on *specific* aspects of Watercare’s operations and long-term performance, and the *specific* strategic responses that the company might contemplate to address the risks. This information was captured and represented in the form of a series of summary templates. Despite this progress, however, those descriptions were still characterised by significant uncertainty, reflecting the fact that many of the risks

involved considerable unknowns.

It was a notable feature of the CRM's work here that he explicitly acknowledged and accounted for this uncertainty in his Strategic Risk Report through what he called the concept of Knowledge Maturity (Clement 2008e, p 17). That is, the CRM's report summarised both the identified strategic risks *and* the state of knowledge on which those descriptions were based (differentiating between *known* risks, *estimated* risks, and *guesstimated* risks). With the Operations and Project Management risk frameworks, as well as in his efforts to promote quantitative risk modelling across the organisation, the CRM ran into certain resistances and dilemmas due to a failure to account for how uncertainty arising from an actor's state of knowledge affects the ability of that actor to specify risk. In this case, the CRM encountered significant uncertainty from the outset, and acknowledged that such uncertainty would always be a feature of any discussion of strategic risk due to the nature and scope of the issues involved. Acknowledging the state of knowledge underpinning descriptions of strategic risks became an integral feature of the strategic risk framework.

### **Performing the role of knowledge facilitator**

A striking feature of the CRM's work on the strategic risk framework is that he appeared to perform considerable cognitive work *on behalf of* Watercare's managers. Not only was he actively guiding them through the risk identification and evaluation process, but in crafting the overarching analytical framework, and then organising and representing the emerging information within that framework, the CRM was, in a very real sense, also "thinking for" Watercare's managers. It can be suggested that the CRM's role in this regard was analogous to that of a *knowledge facilitator* (Roth, Berg, and Styhre 2004).

The concept of knowledge is controversial (Alvesson and Kärreman 2001; Cook and Brown 1999; Snowden 2002; Kreiner 1999; Spender and Scherer 2007; Tsoukas and Vladimirou 2001). The concept overlaps many others, including 'beliefs', 'thoughts', 'perceptions', 'understanding', and 'judgement'; it is somehow different to data and information (Bell 1999); it implies a relationship with 'action' (Buenaño 1999); and, conceived as an object of managerial responsibility, knowledge becomes an asset or resource that firms can draw on to achieve competitive advantage in the knowledge-based economy (Choo 2006; Drucker 1969; Kreiner 1999). For the purpose of justifying the *knowledge facilitator* analogy it will, however, be sufficient to summarise the now well known distinction which describes knowledge as a dialectical interplay between the two ideal types of tacit and explicit knowledge (Nonaka and Takeuchi 1995; Polanyi 1962). Briefly, tacit knowledge may be

understood as personal knowing (Polanyi 1962), as ‘knowing-how’ and therefore integral to skilful action in context (Ryle 1949), but also as knowledge not-yet-articulated (Nonaka and Takeuchi 1995), perhaps even in-articulable, an implicit theory of action (Argyris and Schön 1974, 1978). Tacit knowledge is reflected in the notion that we can know how to do something, and indeed be highly skilled at it, but yet not be able to articulate the rules (theory) that we are following when we act (Ryle 1949). Snowden encapsulated it with the phrase “We can always know more than we can tell, and we will always tell more than we can write down” (2002, p 102). Explicit knowledge, in contrast, is commonly understood as knowledge that has been articulated, codified, and abstracted, taking the form of inscribed concepts, models, hypotheses, metaphors, and analogies (Tsoukas 2003). Tacit and explicit knowledges are sometimes referred to, respectively, as “know-how” and “know-that” knowledges, after Ryle (1949)<sup>11</sup>.

Knowledge facilitation is a strategy for “managing” knowledge in organisational settings (Roth et al. 2004). The concept of knowledge management rests on a rationale similar to that of risk management, that actors’ performances (and hence the performance of the organisation) can be improved through learning, i.e. the creation of new knowledge. The function of the knowledge facilitator role is to create spaces where such learning processes can take place: “...the knowledge facilitator role is explicitly aimed at orchestrating the distribution of knowledge within the firm... to create arenas and opportunities... for sharing knowledge” (Roth et al. 2004, p 200). Roth et al. describe the implementation of a three stage knowledge facilitation process, by two facilitators, the intention of which was to “create a process with tools that could increase knowledge sharing” between different project groups in the target organisation (2004, p 201). The three stages of the process were (Roth et al. 2004 p 201-206):

- 1) **Legitimise and familiarise:** In this case the facilitators were external academics with little knowledge of the company or its projects. This phase consisted of an interview with the project leader, the purpose of which was to familiarise the facilitators with the project and for the facilitators and the project leader to agree on how to proceed.

---

<sup>11</sup> The concepts of tacit and explicit knowledge are idealisms. While it is possible, at a very broad level, to conceptualise one form of knowledge as that which resides “in someone’s head” and another as that contained in reports, text books, encyclopedias, presentations, etc., at any deeper level the concepts intertwine such that distinctions between them, if they can be made at all, become extraordinarily subtle. Indeed, even the notion of “explicit knowledge” can be interpreted as the product of introspective reflection, and therefore not necessarily involving a physical process of inscription “outside one’s head”.

- 2) **Unlock tacit knowledge – Structure explicit knowledge:** This step consisted of two brainstorming sessions with the project team. Information from the first session, collected on flipcharts and whiteboards, was subsequently collated and structured by the facilitators, and then discussed in a second session. The “time between the sessions gives the project members a chance to reflect on the material, and this often leads to more experiences from the first session being added. The facilitators lead the sessions and structure information, whereas the project members hold an interactive dialogue. This results in a mind map, with the common knowledge gained from all of the sessions...” (p 203).
- 3) **Share with the organisation:** This step was an interactive seminar where the results of the facilitation process were shared (fed back) to other members of organisation. The seminar involved both a traditional presentation format as well as open dialogue between project team members.

Through the above process, the project team members were encouraged to articulate their tacit knowledge of project experiences (externalisation), and to share that knowledge via interactive social processes (socialisation) (Nonaka and Konno 1998; Nonaka and Takeuchi 1995). The facilitators guided the process but also contributed to the creation of knowledge by structuring the information which emerged from the brainstorming sessions (combination). The above process performed by Roth et al.’s knowledge facilitators was very similar to the process facilitated by the CRM with Watercare’s senior management. The same three stages are evident:

- 1) **Legitimise and familiarise:** Unlike Roth et al.’s facilitators, Watercare’s CRM was an internal actor within the company and was therefore familiar, to a certain extent, with the organisation and its context, projects, and processes. Indeed, the level of detail to which the CRM was able to pursue the restructuring (next stage, below) and “fleshing out” of the identified strategic risks was only possible due to his personal knowledge of the Watercare organisation. Legitimacy was conferred through his role as Corporate Risk Manager, i.e. giving him the necessary authority to claim responsibility for guiding the strategic risk assessment.
- 2) **Unlock tacit knowledge – structure explicit knowledge:** In this case, the brainstorming was not carried out in a group session, but rather by the participants individually using the templates provided by the CRM. The resulting information was subsequently restructured by the CRM, as described above, and then fed back to the

participants in a group session. As with the process described by Roth et al., this provided the participants an opportunity to reflect on the material and to question what had been presented or add additional information as appropriate. Rather than a mind map, the end product in this case was a structured report, authored by the CRM, which documented the process, the strategic risk framework, and the agreed details of the identified strategic risks.

- 3) **Share with the organisation:** in this case, distribution of the report to the participant managers, and subsequent discussion at strategic planning seminar.

On the basis of the above comparison, there are strong parallels between the work performed, respectively, by Watercare's CRM and the knowledge facilitators described by Roth et al. (2004). Using various tools (e.g. the PESTEL template, the Strategic Risk Report and accompanying risk templates, PowerPoint© presentations, etc.) the CRM created a process and spaces in which Watercare's managers could make explicit their tacit knowledge about the company's strategic risks. Although the results of the initial brainstorming were messy, by collating, rationalising, organising, and formatting the managers' initially vague responses into a comprehensible and coherent framework, the CRM enabled them to reflect on and think through the issues in more specific detail, and, importantly, to make that specification explicit. The CRM contributed significantly to this process by thinking through the overarching analytical framework and combining the available information to produce "new knowledge" of Watercare's strategic risks.

Drawing on von Krogh, Ichijo and Nonaka's (2000) description of a knowledge activist, Roth et al. (2004, p 206-207) described the knowledge facilitators as fulfilling several roles: serving as a catalyst and co-ordinator of knowledge creation initiatives, guiding the overall direction of knowledge creation activities, facilitating trust between organisation members, and enabling a sharing culture. While I do not have direct evidence that the CRM was able to improve the level of trust or enable a proactive sharing culture amongst Watercare's managers, it is clear from the earlier description of his interventions that the CRM acted as both a catalyst and co-ordinator of knowledge creation activities, and, through his efforts to structure the strategic risk framework, provided the overall direction for the process. In contrast with the CRM's other endeavours, here he apparently succeeded both in developing the strategic risk framework, and improving the quality of Watercare's strategic "risk data". The end product of the facilitative process described above was the Strategic Risk Report which documented the process, explained the overarching strategic framework, and collated the extant knowledge about Watercare's strategic risks in a comprehensible format



according to a rigorously defined analytical framework. While the report was by no means final and perfect, it did embody considerable evolved understanding about Watercare's strategic risks, and constituted a significant improvement in data quality over the original PESTEL template (see Table 4.2). Prior to the CRM's intervention, no such comprehensive or coherent repository of strategic "risk data" existed.

*Table 4.2. Evolution of the strategic risk framework and template*

	<b>Original PESTEL framework</b>	<b>CRM's redefined framework</b>
Framework and template design	<p>Generic definition of "strategic" as significant, long-term issues.</p> <p>Common framework in strategic management literature. Six generic categories for classifying perceived trends in the external business environment</p> <p>Simple matrix format, no definition in framework (e.g. sub-categories, parameters etc.)</p>	<p>Clear, unambiguous definition of "strategic risk" for Watercare's context.</p> <p>Bespoke risk template tailored to Watercare's specific operating context.</p> <p>High level of detail and coherent structure organised around progression from cause to consequence to response.</p>
Data quality	<p>330 "strategic issues" identified.</p> <p>Vague, ambiguous descriptions with little detail (often just a few words).</p> <p>Little structure. Effectively the results of an unstructured brainstorming session.</p>	<p>23 strategic risks defined and described.</p> <p>Clear, unambiguous descriptions, moderate level of detail.</p> <p>High degree of structure linking internal and external trends to impacts on enterprise systems and performance, and potential strategic responses.</p>

## Discussion

The core of the Corporate Risk Manager's vision for developing Watercare's risk management capabilities was a number of proposals to improve the quality of the company's "risk data". The CRM's primary goal for these interventions was the production of good "risk data". That is, he expected that more objective risk frameworks and registers, rigorous analytical methods, and quantitative modelling capabilities would lead to the production of "risk data" which objectively represented the systems, processes, and environments about which and within which decisions were made. This goal was rooted in his assumed decision support rationale, i.e. that there was a positive correlation between the quality of organisational decision making and the quality of the information on which those decisions were made (see

Chapter 3). The CRM was, however, relatively unsuccessful in achieving this goal through the endeavours reported in this chapter (see Table 4.3).

*Table 4.3. The CRM's score card*

<b>Task objectives</b>	<b>Success relative to objectives</b>	<b>Impact on data quality</b>
<i>Task: Redefining the Operations Risk Framework</i>		
To constrain the subjectivity of the risk assessment and communication process by specifying the risk framework as an objective representation of the corporate objectives and operational performance standards and parameters.	<b>Achieved:</b> Produced a revised risk assessment framework which was justifiably “more objective” than the company's existing framework.	<b>Minor:</b> The new framework was difficult, time consuming, and costly to produce and maintain, and was more complex than the existing framework (which would make it difficult for staff to accept as an effective management tool).
<i>Task: Redefining the Project Risk Register</i>		
To tailor the project risk register to accurately record “risk data” in terms particular to the project management context.	<b>Achieved:</b> Redefined the risk register to more accurately reflect the objectives and performance requirements of the project management context.	<b>Minor:</b> The new register was difficult to use because it was not well suited to recording ambiguous or uncertain information.
<i>Task: Facilitating the identification and assessment of strategic risks</i>		
To collate and organise available knowledge of Watercare's strategic risks within a coherent framework.	<b>Achieved:</b> Produced the Strategic Risk Report which explained the overarching strategic risk framework, and described the identified strategic risks in the form of a series of structured templates.	<b>Significant:</b> The Report was a comprehensive collation of extant knowledge of the company's strategic risks within a rigorously defined analytical framework, constituting a new knowledge repository for the organisation.

The idea of extending the company's reliability modelling capabilities quickly encountered the fundamental limitation of data availability and quality, and thus didn't really “get off the drawing board”. Of the other three tasks reported here, the CRM was successful at least with respect to his initial objective for each task. That is, he did produce a more detailed and objective risk assessment framework and corresponding risk registers, and he was successful in collating information about Watercare's strategic risks. But only the latter task had any discernable impact on the quality of the company's “risk data”, in the form of the Strategic Risk Report which collated extant knowledge of the company's strategic risks in a comprehensible format according to a rigorously defined analytical framework. While the report did not contain an exhaustive analysis of each risk, prior to the CRM's intervention no such comprehensive or coherent repository of strategic “risk data” existed. In this sense the

report represented a significant improvement in the quality of Watercare's strategic "risk data". In contrast, the new risk framework and risk registers were found to be problematic in use. The new risk framework was perceived to be too complex, while the detailed risk register was ill-suited to capturing ambiguous or uncertain information, limiting its useful employment to those rare situations where risks were already clearly and precisely defined.

The following sections engage in a metatheoretical discussion over the ontological and epistemological status of risk in order to frame risk as a product of inquiry. That conceptual framework is then used to frame up lessons learned from the CRM's experiences about designing ERM infrastructure.

## **Ontological lessons**

The CRM's primary concern about Watercare's existing risk framework was that it allowed considerable subjectivity in risk assessment and communication, a fault which the CRM felt was attributable to the fact that the framework was "one-size-fits-all". The one framework was used for assessing risk right across the enterprise, from project planning, design, and management, to operations, to strategic planning. For the CRM this was problematic because a single framework could not adequately represent the specifics of each context. He therefore sought to develop three frameworks tailored specifically to the different operational contexts of the organisation: Project Risk, Operations Risk, Strategic Risk. His approach was to limit the subjectivity of the risk framework in use by seeking clarity and detail in the design of the framework. He envisaged that the new risk framework should be, as far as possible, an objective representation of the company's objectives and performance standards.

However, the CRM discovered, somewhat ironically, that there was an inherent subjectivity to all the decisions involved in the design of the risk framework: decisions about which objectives to represent, which performance parameters or criteria to use, about the relative priorities of those objectives and parameters, and about how to organise them into a coherent framework. Contrary to the CRM's initial assumption, these decisions could not be reached via a logical deconstruction of the corporate objectives and performance standards contained in various texts within the organisation. This subjectivity pointed up the question of the ontological status of risk. That is, is risk fundamentally a real object which exists independent of what humans may think about the matter, or fundamentally a social and cultural object, arising out of social discourse and therefore only existent within that discourse (Althaus 2005; Jaeger et al. 2001; Renn 1992; Rosa 1998; Schrader-Frechette 1991)?

The common, everyday notion of risk is that associated with the chance of loss or adverse

consequences arising from some peril or hazard, which may be environmental or the result of human activities (Collins English Dictionary 2000; Fischhoff et al. 1984; Hilgartner 1992; Kates and Kasperson 1983; Merriam-Webster Online 2008; Short 1992; Ward 2005). This definition reveals the conceptual category of risk to be the product of two essential criteria: chance (possibility) and loss (harm); more commonly referred to as the dimensions of Likelihood and Consequence in Risk Management terminology.

In the field of corporate risk management, however, it has more recently become normal to talk about risk as encompassing positive as well as negative outcomes. Rosa (1998), for example, argues that the notion of risk refers to situations where something of human value is at stake, regardless of whether the anticipated outcomes are positive or negative. Renn (1998, p 51) employs a similar definition, while Ward (2005, p 39) lists ten post-1997 definitions of risk from various industry sources all of which define risk in terms of impacts on outcomes or objectives, whether positive or negative.

The concept of risk has been appropriated by a wide range of disciplines across the physical and social sciences, such that it is possible to talk about a taxonomy of perspectives on risk (Althaus 2005; Renn 1992). The various disciplinary perspectives reflect quite different and often opposing ontological and epistemological orientations. The realist/positivist paradigm sees risk as an objective phenomenon residing in the real world that can be identified, quantified, and understood as fact via logico-scientific methods independent of human perceptions, values, and politics (Rosa 1998). This paradigm forms the basis for technical risk analyses in the physical, biological, and natural sciences, in engineering and technology, and in medicine (Althaus 2005). Theoretical discourse in these domains treats risk as something that is specifiable, measurable, and quantifiable (Althaus 2005). Outcomes that matter are assumed to be common sense (e.g. utility, wealth, physical wellbeing, life, etc.) and loss is simply represented by a negative change in the quantity of interest. In this view, the elements of chance and negativity are seen as real states of the world and risk is “at bottom” ontological (Rescher 1983). That is, humans are exposed to risk whether they know about it or not; as in the distinction between taking a risk and running a risk (Rescher 1983).

At Watercare, the dominant notion of risk was that associated with the potential failure of physical assets within the water and wastewater infrastructure networks; which reflects the fact that developing, operating, and maintaining infrastructure assets was Watercare's core business. As a bulk supplier, Watercare's formal market consisted of only six customers, the Local Network Operators, and that relationship was legally defined at certain physical points within the network by a mixture of regulatory and contractual specifications. But the infrastructure systems operated by each company were contiguous, forming an

uninterrupted network across the region. In this sense, Watercare's relationship with the LNOs, and indeed with the people and businesses of Auckland who were the final consumers of Watercare's services, was also clearly physical. What happened in Watercare's network affected what happened in the LNO networks, and ultimately also the end consumer. Although changes in the performance of Watercare's network were generally not of sufficient magnitude to be noticeable to the end consumer under normal operating conditions (i.e. minor daily fluctuations in pressure and quality), this context required the company to maintain an extraordinarily high degree of reliability. Operational failures could have very real and very dire consequences far beyond the impacts to the company's reputation or financial statements. For example, a drought, earthquake or volcanic eruption, the collapse of a dam, equipment and process failure, damage to pipelines by third parties, and incidents in confined spaces, to name a few, were all easily imagined and understood as potential events which, depending on the situation, could result in serious consequences, not just for the company but also for Auckland's population, economy, and environment. Such consequences could include contamination or loss of the water supply, flooding of a sensitive ecosystem or popular beach with sewerage, illness, injury or death of employees or members of the public, damage to public infrastructure or private property, and in addition to the immediate and consequential financial costs, the prosecution of the company and employees, and loss of confidence in the company by its stakeholders. Thus, in Watercare's context, risk was very much a real phenomenon.

At the other end of the spectrum, constructivist approaches to risk research (particularly in psychology, sociology, and anthropology) have highlighted the shortcomings of the realist/positivist paradigm in taking for granted the subjective value judgements inherent to the notion of loss or harm (Althaus 2005). This paradigm argues for the significance of psychological, organisational, political, and cultural factors on perceptions of and attitudes toward risk (in other words that context matters, Renn 1992). In this view, risk is fundamentally a subjective and therefore epistemological category (Thompson 1986). Not only is the conjoining of chance and negativity into a single concept a distinctly human endeavour, but, viewed historically, it is also a relatively recent idea representing an epistemological revolution from notions of a world ruled by fate to a world ruled by probability (Bernstein 1996). The element of negativity is not solely an objective fact of the world, but rather is the result of a human value judgement, it "becomes available as an element of risk only on the condition that a state of affairs has been judged, perceived or conceived to be bad, at least relative to some other state of affairs" (Thompson 1986, p 279).

Although risk in Watercare's context was associated with very real consequences, both

for the company and for others, various stakeholders expressed quite different and apparently subjective perceptions of what was significant, and why. The CRM's discussion with Operations line managers (Dialogue 4), for instance, revealed the complex range of inter-related standards to which the company was required to perform. Although the managers who participated in that discussion were all knowledgeable and experienced individuals, and the group made a genuine effort to arrive at an objective prioritisation of outcomes, they ultimately failed to agree on the relative significance of different performance failures. It was clear from the discussion that each had a personal set of criteria for evaluating the significance of outcomes which was influenced not only by their specific roles and responsibilities in the company, but also by their own subjective expectations of who might be held accountable in the event that things went wrong.

Similarly a discussion with Watercare's General Managers revealed that they did not perceive the objectives in the company's Statement of Corporate Intent (SCI) to be important with respect to the strategic development of the water and wastewater infrastructure. This was a surprise to the CRM because the SCI was ostensibly the formal mechanism through which Watercare's Board and Shareholders could specify how the company was to give effect to and be judged against the firm's principal objectives set forth in the Local Government Act and the Company Constitution. On the contrary, Watercare's General Manager's saw the SCI as merely an aspirational document which described what sort of organisation the company should aspire to be and how it should operate vis-à-vis the delivery of its services, but not as a fundamental driver of the business. For Watercare's General Managers, strategic investment in water and wastewater infrastructure was driven by three primary forces: regional growth in water demand/wastewater volumes, changing regulatory requirements, and improvements to levels of service (usually at the request of the Local Network Operators).

Although the realist and constructivist perspectives are often presented as lying at the extreme opposite ends of a continuum of ontological positions, both are correct with respect to their claims about the ontological status of risk. As a label risk does refer to real states of the world, but, at the same time, what counts as significant, and therefore as risky, is the product of social value judgements. Some state of the world must be recognised and judged by someone in order for risk to actually exist. Rosa (1998, p 23) argued that this duality makes it easy to conflate the ontology of risk (i.e. questions of "what is the nature of the world of risk?") with the epistemology of risk (i.e. questions of "how do we understand and know risks?"). To avoid the resulting "theoretical and logical mischief", Rosa (1998, p 23) suggested that we must distinguish between the notion of risk on the one hand, and the real states of the world to which that label refers, on the other. Specifically, any statement of risk expresses

expectations about three things: (i) a future state of the world resulting in certain outcomes, (ii) how significant (good or bad) those outcomes will be, and (iii) how likely or unlikely those outcomes may be within a specified period of time. Rosa's point is that for such a statement to be sensibly uttered and understood, the speaker must intend, and the recipient must assume, that the statement represents: (a) certain aspects of the real world (elements i & iii, above), as well as (b) what some person or group actually believes about the significance of those outcomes (element ii above). In this regard, risk can only be said to exist when (Rosa 1998):

- Something is *actually* at stake (there must be something real to be lost or gained); *and*
- What is at stake is judged to be of value to someone (someone must be concerned about the outcomes).

The latter component makes risk a contestable phenomenon since humans can define what is significant based on any criteria of their choice. That choice is always subjective to the extent that it depends on personal perceptions of “what matters” in any given context. This is not to say, however, that risk is an arbitrary or capricious category, or that every perspective has the same standing as every other. Risk is not just a matter of personal whim. Rather, perceptions of risk are shaped by various social, cultural, political, and technological factors, and what those factors are and their relative importance depends on one's positionality (Hilgartner 1992; Rosa 1998, p 28). Thus, in any particular context there may be multiple legitimate conceptualisations of risk, depending on the number of stakeholders and their relative priorities.

It was for this reason that the CRM could not define an “objective” risk framework simply by pursuing a detached deconstruction of the statements of corporate objectives and performance standards contained in various texts within the organisation. Any such deconstruction would imply a certain prioritisation of outcomes and the CRM discovered that his initial prioritisations did not always agree with those of organisational stakeholders. Indeed, the CRM discovered that different people prioritised outcomes quite differently, even where it seemed that certain outcomes were “obviously” more important than others (for instance that “water” was more important than “wastewater”; Dialogue 4). Because of this, the CRM's initial attempts at specifying the risk framework were strongly contested by various stakeholders in the organisation as not accurately reflecting what really mattered.

In order for the CRM to specify a framework which did accurately represent priorities throughout the organisation, he had to uncover and make explicit the perceptions of what mattered to various stakeholder groups in their specific functional contexts (i.e. general

managers in the strategic context, project managers in the project delivery context, and operations managers in the operations context). These were quite specific understandings about how things worked in each context, and about what was important and what was not, and which were not readily apparent to the CRM as an outsider. Such understandings were acquired by the various actors through long practical experience, both in their professional fields in general, and in their specific roles within the company, and were only accessible to the CRM via frequent and extensive engagement with those actors in those contexts. The CRM's experience thus reflects the imperative commonly expressed in world-level frameworks and standards that Risk Management tools need to be appropriate to the contexts in which they will be used, but it also reveals the task of defining those contexts to be, as the CRM described it in Dialogue 14, "more art than science".

The process was also one of tracing connections between the corporate objectives and the business systems, processes, tools, and capabilities designed to fulfil those objectives at various levels in the organisation. In particular, the CRM was able to establish links between ideas of "what mattered" in each context and the existing capabilities for data capture, analysis, and reporting within the company. Since the existence of those systems and metrics was itself evidence that certain performance parameters were important to the company, by specifying those parameters in the risk framework, the CRM effectively enrolled (Latour 1987, 1988) those existing calculative infrastructures in support of his new framework. The connections that the CRM traced through his engagements with organisational stakeholders thus served to support and justify his framework as a "more objective" representation of the business than alternative specifications (specifically, as "more objective" than the company's existing risk framework).

## **Epistemological lessons**

Within a given system there may be a wide range of sources of risks, including plant and equipment failure, variation of input parameters, human error, or impacts from external events. A central principle of Enterprise Risk Management is the notion of an holistic, cross-functional approach, in contrast to the traditional treatment of risks in functional "silos" (Lam 2003; Power 2007; Ward 2005). Such an approach seeks to address interdependencies between risks at the enterprise level that would otherwise be missed by traditional approaches to risk management. A central concept is that of the Enterprise Risk Profile, conceived in abstract terms as an aggregate representation of the firm's risk universe, accounting for all types of risk exposures across different business silos and for the inter-



relationships between those exposures (see Appendix V).

The Enterprise Risk Profile concept assumes that it is possible to calculate a probability distribution of outcomes attributable to each discrete source of risk. These separate distributions may then be aggregated by probabilistic modelling (e.g. Monte Carlo simulation) to produce an aggregate probability distribution for the variation in performance of the system as a whole, accounting for all identified sources of risk. The CRM proposed that Watercare's enterprise risk profile should be made calculable via such quantitative risk modelling methods and to this end proposed to extend the company's existing reliability modelling capabilities. But in an organisation dominated by engineers and engineering work there is a certain irony to the fact that the CRM's vision of extending those capabilities was thwarted by a scarcity of hard data. This highlighted important limitations with respect to when and where probabilistic models can meaningfully be employed to aggregate individual risk exposures to the enterprise level; limitations which stem from the propagation of uncertainty through such models.

Similarly, there is a certain irony in the fact that the CRM succeeded in redefining Watercare's project risk framework and register to more objectively represent the project management context, but in doing so produced a tool that was difficult to use for recording "risk data" during a project risk identification workshop. That particular difficulty arose because elements of the project design were still uncertain and the workshop participants lacked detailed knowledge of the relevant risks. This experience called attention to the fact that risks can only be fully specified toward the end of the risk assessment process, once all of the investigation and analysis has been completed. Earlier in the process the relevant questions have not been answered in full, and ambiguity and uncertainty may remain about objectives, system components, their functions and relationships, and about external variables. Detailed risk registers are relatively incompatible with such contexts.

The common feature of both of the above cases was a discrepancy between the normative definition of risk, as implied by the CRM's detailed risk register (and the risk profile concept), and the information and knowledge which was actually available in certain contexts and on which basis risks could be specified. This discrepancy pointed up the question of the epistemological status of risk. That is, what is the difference between Risk and Uncertainty?

The distinction between Risk and Uncertainty can easily be confused, not least because risks are, by definition, future outcomes about which there is uncertainty. The second dimension of risk is that of chance or possibility, which means that the expected consequence (i.e. loss or gain) is only more or less likely, but not certain, to occur or not to occur over some future time period (Rosa 1998). Renn (1998) stated that if outcomes are

guaranteed then the term “risk” makes no sense because we know precisely what will be lost or gained and when (a situation usually referred to as fate). This is conceptually represented by the left-hand side of Figure 4.2 where risk lies between the limits of the likelihood axis, i.e. where the probability of outcomes is  $0 > P(x) < 1$ .

Strictly speaking, the correct mathematical representation of risk is as a probability distribution which maps the range of possible outcomes with their corresponding likelihoods. Traditionally, risk has been distinguished from uncertainty by the quantification of that probability distribution. The seminal distinction was made by Knight (1921), who defined three probability situations (Runde 1998, pp 540-543):

- 1) Where it is possible to assign a numerical probability *a priori*, on the basis of general principles, to outcomes taken to be equally probable and mutually exclusive except for really indeterminate factors – such as in assigning the numerical probability of 1/6 to the chance of rolling a ‘4’ with a standard die;
- 2) Where it is possible to derive numerical probabilities *a posteriori* via an empirical classification of outcomes obtained in classes of more or less homogenous trials – such as using historical data to calculate the chance of dying in a motor vehicle accident across a given population and geographical area; and
- 3) Where there is no valid basis of any kind for classifying instances – we simply do not know enough to assign a numerical probability.

Knight (1921) referred to the first two as cases of risk (i.e. probabilities are available), and to the latter as uncertainty (i.e. information is too imprecise to be summarised by probabilities); a distinction which has become ingrained in economic and decision theory (Runde 1998). In Knight’s terms risk is not just a claim about certain things in the world, but is also a state of knowledge about those things. That is, we can talk about risk when we know enough to reasonably calculate and specify the real probabilities of outcomes; where we cannot make that estimate then we are merely uncertain (Knight 1921; Runde 1998).

However, while the above categories are conceptually simple, and indeed normatively attractive, such a sharp distinction between risk and uncertainty does not exist in practice, or even, for that matter, in theory. Risk assessment literature, for instance, often distinguishes between *aleatory* uncertainty versus *epistemic* uncertainty (Hora 1996). The former arises from the natural, unpredictable variation in the performance of the system in question, and is assumed to be irreducible but quantifiable, while the latter type of uncertainty arises from an actor’s lack of knowledge about the system, and, it is assumed, can be reduced through

further inquiry (Hora 1996). The purpose of a risk assessment is to quantify the aleatory uncertainty on the performance of a system; this being an estimate of risk in Knight's terms, above. But since human knowledge is always limited to some extent, any such estimate will be imperfect and therefore more or less uncertain in the epistemic sense. This epistemic uncertainty is fundamentally ineradicable<sup>12</sup>, and can only be reduced through further inquiry to a point where it may be considered negligible for the purposes of the calculation at hand (what Knight referred to as "really indeterminate factors").

Runde (1998) suggested, in this regard, that Knight's trichotomy, above, should be seen not as three discrete categories, but rather as representing a continuum of situations where an actor's knowledge of the probability of events moves from less certain to more certain. In this sense, the initially sharp theoretical distinction between risk as quantitative knowledge and uncertainty as qualitative knowledge is replaced by the inherently fuzzier distinction of more versus less support for any particular estimate of probability. This is illustrated in Figure 4.2, overpage. The left-hand side of the diagram portrays the notional "risk space" which represents the three essential components of risk: (i) some real future state of the world resulting in certain outcomes, (ii) how significant (good or bad) those outcomes will be, and (iii) how likely or unlikely those outcomes may be within a specified period of time. The right-hand side of the diagram fuzzily distinguishes risk from uncertainty on the basis of what some actor might actually know about (i), (ii), and (iii).

In practice an actor's state of knowledge about some risk situation advances from Uncertainty to Risk (i.e. bottom-left to top-right of the right-hand figure above) through some process of inquiry. In any given situation actors may be uncertain about what the relevant risks are, and may only be able to talk about them in vague and ambiguous terms. But by the time they have properly defined, analysed, and evaluated those risks, they should be certain about the relevant risks and be able to talk about them in clear, specific, and quantitative terms. Only in the latter case, however, where actors can calculate and specify the probability of future outcomes with objective precision, is their use of the term "risk" consistent with the abstract theoretical definition of the term. The CRM's relative successes and failures with his various interventions discussed in this chapter may be explained, at least in part, in terms of the timing of his interventions with respect to this knowledge production context (this is illustrated in Figure 4.3, overpage).

---

<sup>12</sup> Human inquiry faces a number of practical obstacles in the pursuit of true knowledge, such that we can never establish an absolute correspondence between the reality of the world around us and our knowledge of that world (Loasby 1999; Megill 1994; Rescher 1995, 2005).

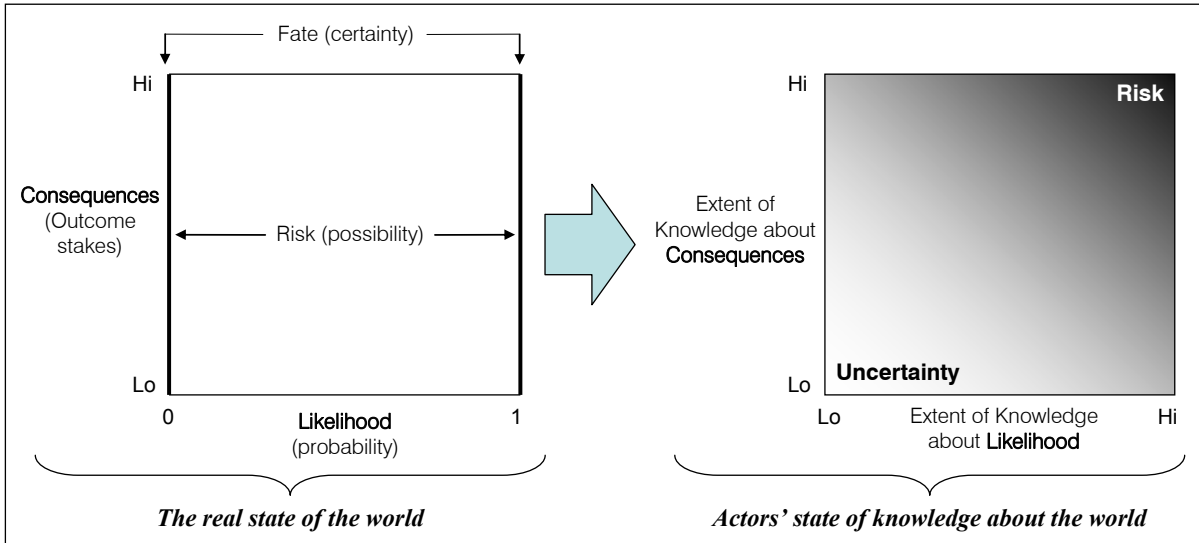


Figure 4.2. The distinction between Risk and Uncertainty (modified from Rosa 1998, Figure 3, p 35)

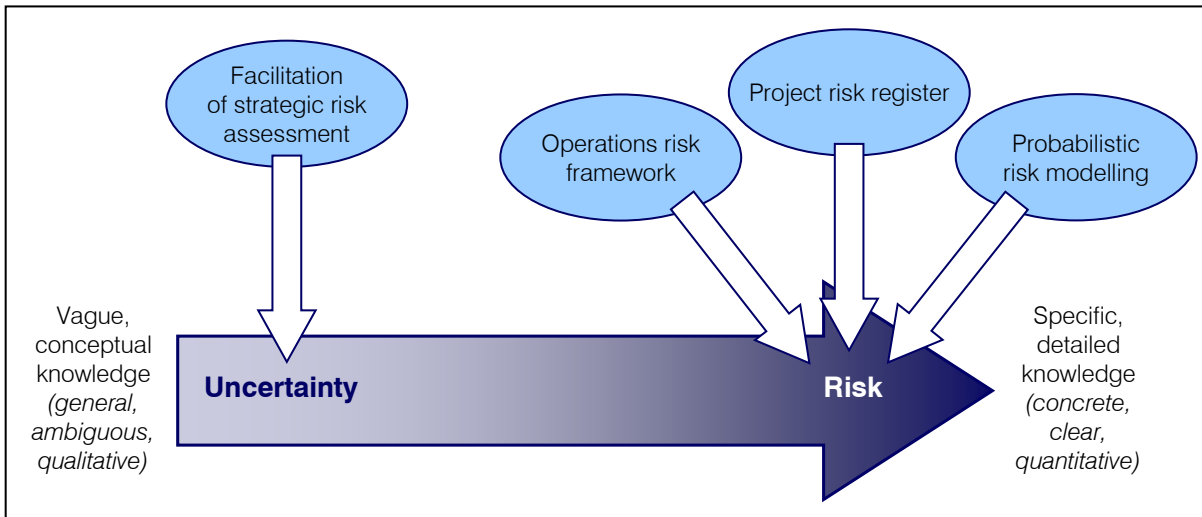


Figure 4.3. The assumed knowledge production contexts of the CRM's interventions discussed in this chapter. The dark shaded arrow represents states of knowledge about the world as a progression from uncertainty to risk (cf. Figure 4.2)

I had difficulty filling out the CRM’s new project risk register in a project risk identification workshop because the new register was a rigorously defined and detailed template for recording data about risks (in the latter sense, above), but the participants in the workshop were discussing mere uncertainties (in the former sense, above). The register was effectively incompatible with the context in which it had been employed. From this observation, several important limitations were identified with respect to the productive use of detailed risk registers:

- To the extent that the detailed risk register can be usefully employed only toward the

end of the risk assessment process, it does not support that process. That is, the risk register does not help actors to produce knowledge about risks, but simply records that information once it has been produced.

- As a standardised template for formatting and recording (inscribing) the outputs of the risk assessment process (i.e. statements of risk), the risk register primarily fulfils a bureaucratic function. The register is effectively a defensible record of the risk assessment and decision making processes, and serves as a baseline for monitoring, internal control, and audit of risk controls.
- While not denying that a clear and precise record of risks is valuable to the organisation, particularly as a basis for future communication, that value is only realisable when the data in the register is supported by objective knowledge of the risks in question. If actors do not possess such knowledge then either they will be unable to express the risk statements in the terms of the register (see next point), or, if they do, those statements will appear to be more objective than they really are. Depending on the format and detail of the register template, this can have the effect of masking uncertainty, subjectivity, and ambiguity in actors' actual understanding of those risks.
- In order to fully specify risks within a detailed register template, with clear differentiation between causes and consequences, quantified measures of consequence, and specific controls clearly addressing either likelihood or consequences, requires important distinctions to be made. That is, a more detailed register requires users to both possess and express a deeper and more refined knowledge of the risks in question. Since those distinctions only become possible toward the end of what might be a fairly lengthy and involved process of inquiry, the risk register can be seen as imposing a uniform standard level of detail on the outputs of the risk assessment process. In principle, all risks entered into the register should be defined to the same level of detail, which translates into a greater (and more costly) burden-of-proof than might otherwise be required.

Furthermore, since the risk register simply reflects the terms of the risk assessment framework, the above limitations also apply to the framework itself. A similar conclusion was reached with respect to the CRM's promotion of quantitative risk modelling. Mathematically, such models amplify or compound the uncertainties on individual distributions, such that the uncertainty on the resulting combined distribution may be much larger than for the

input distributions separately (Dai et al. 2007; Yin et al. 2001). Thus, in order for aggregate risk profiles to be mathematically meaningful, the uncertainty on each individual risk distribution input into the model must be very small; which implies: (i) that probabilistic risk models should only be used to aggregate risks that have been quantified with reasonable certainty; and (ii) that probabilistic risk modelling is costly because considerable resources have to be devoted to the analysis of individual risks to ensure such certainty. The CRM's vision of extending Watercare's risk modelling capabilities was thwarted in this regard by the relative absence of comprehensive and reliable asset condition information for infrastructure beyond the company's water and wastewater treatment plants.

The knowledge production context also explains the CRM's relative success with the strategic risk work. In that case, the framework and register were developed as part of a facilitative process where the CRM guided Watercare's senior managers through an explicit and systematic exercise of strategic context definition, risk identification, and assessment, spread over several workshops. The process was an iterative one of extracting information from Watercare's managers, and collating that information into a comprehensible format, with each step informing the subsequent work. The earlier description of that process covered one iteration: the initial collection of information using the common PESTEL template, the collation of that information and our attempts to make sense of it within a coherent definition of what "strategic risk" meant in Watercare's context, and the subsequent re-engagement with Watercare's managers using the new framework and template to further refine the risk definitions. The strategic risk framework and register were therefore products of the process they were intended to support rather than pre-defined devices brought "ready-made" to the table.

The important insight is that the level of detail in the strategic risk framework evolved to suit the needs of the knowledge production process. In the beginning the process was characterised by uncertainty, ambiguity, and vagueness, and the PESTEL template was appropriate as a framework for capturing information in this context. It offered a simple set of categories for classifying factors, issues, and trends in the external environment, while also being sufficiently generic as to not constrain or bias the brainstorming process. This is not to say that the CRM chose to use the PESTEL framework for these reasons; he didn't. Rather, at the time, the CRM was himself uncertain about what "strategic" meant in Watercare's context, and the PESTEL template merely provided a convenient and well-understood framework for thinking about the external environment of the organisation. One could suggest that the framework itself contributed to the vagueness and ambiguity of the data that was eventually collected; a more detailed and structured framework might have

elicited better responses from Watercare's managers. This may be true, but the point is that the PESTEL template served as a quick and easy means to elicit a broad picture of the range of issues and trends in Watercare's environment that were perceived to be important by the company's managers.

The subsequent work of collating and rationalising that initial data set (of 330 issues) revealed the limitations of the PESTEL template. The initial set of issues could be rationalised to a certain degree within the PESTEL framework, but ultimately this was only going to produce a smaller list of issues still classified under or across the same six categories. In order to communicate the rationalised list of issues back to Watercare's managers, and to subsequently elicit more detailed information about the causes of those issues, their impacts on Watercare's business, and the potential responses the company might pursue, a more detailed framework, structured to this task (i.e. to Watercare's specific context), was required. This eventually took the form of the strategic risk definition and an earlier version of the template described on page 116. Compiled as a report and supported by a PowerPoint presentation, the new framework and template served as a basis for the second round of engagement with Watercare's managers. The additional information collected from that process was then subsequently used to further refine the framework. In this way the strategic risk framework and template evolved in tandem with our understanding of the identified strategic risks.

### **The value problem: how much detail?**

The CRM was well aware that risk management infrastructure and tools should be appropriate to the contexts in which they will be employed. This imperative underpinned his concerns about the lack of context-specificity in Watercare's existing (2003) risk framework. The CRM assumed that in order for the risk framework to efficiently and objectively moderate communication between different communities of practice (i.e. different audiences) in the organisation, it would need to effectively marry together the terms and parameters with which those audiences were familiar. This suggested that different frameworks would be required for different functional contexts in the organisation to achieve the desired degree of objectivity. For each framework the structure (consequence categories) would be provided by a deconstruction of the corporate objectives, while the content (evaluation criteria and parameters) would be provided by the various performance standards relevant to that functional context. This "marrying together" would allow a risk framework to be tailored specifically to the project management context, for instance, in

terms familiar to the project managers (the “actors at the coal face”), while ensuring that the resulting risk assessments would be represented in terms familiar to management (the “executive audience”). After considerable effort the CRM did produce a series of new risk frameworks which he claimed were “more objective” than Watercare’s existing “one-size-fits-all” framework. Although that claim was, in and of itself, justifiable, the more detailed frameworks also posed two important dilemmas:

- They were difficult, time consuming, and costly to produce, and for this reason were foreseeably inflexible in the face of changing objectives and performance standards; and
- They constituted a more complex and potentially more cumbersome metrological system which would make it difficult for staff to accept them as effective management tools.

The CRM’s interventions were also grounded in an assumption that risk assessments did, or should, produce “risk data” with clear differentiation between causes and consequences, quantified measures of consequence, and specific controls clearly addressing either likelihood or consequences. This assumption was underpinned by the CRM’s overarching decision support rationale that good decision making should be based on good data (Chapter 3). However, the implementation of detailed risk registers and comprehensive risk models encountered two critical obstacles in practice:

- Such detailed infrastructure is not productive in the early stages of inquiry which are marked by uncertainty, ambiguity, and vagueness about both means and ends; and
- Not all inquiries produce such detailed “risk data”. Indeed, it may be noted that, while it is a general rule of thumb that the more significant the perceived risk the greater the resources given to understanding that risk, this does not necessarily imply that a firm’s biggest risks are the best understood. The most significant risks may involve large-scale, long-term changes in a firm’s strategic context, which may not be amenable to quantification with reasonable certainty.

This points up that, in addition to being appropriate to the functional contexts of the organisation, ERM infrastructure should also be appropriate to the knowledge production contexts in which it will be employed, if it is to be used effectively. So, for instance, detailed risk frameworks are of little use in the project definition phase due to the uncertainties involved, while, vice versa, simple, generic frameworks are of little use during the final stages



of construction planning where there is a need to collate and organise a considerable amount of detailed information about risks and risk controls. The practical problem for CROs is to understand the different needs of actors engaged in different phases of the knowledge production process, and which tools will be appropriate to support those needs.

More broadly, however, the lessons from the CRM's experiences call attention to the fact that the normative goal of objectivity must always encounter the inevitable practical constraints to its achievement. The overarching value problematic can be framed as follows. Taken to its ultimate extreme, the CRM's goal of an objective risk framework would amount to representing all of the organisation's objectives, performance standards and parameters in absolute detail, and in a format that made explicit all of the relative priorities between those objectives and standards. In other words, the risk framework would have a 1-to-1 correspondence with the company's metrological performance system. Of course, given the complexity of that system for most organisations, such a framework would be utterly impractical both from a design perspective and in use. The other end of the spectrum is equally undesirable, where the correspondence between the risk framework and the company's objectives and performance standards would be so loose as to promote rather than constrain subjectivity in use. Conceptually, a suitable risk framework must lie somewhere between these two extremes. Practically, however, the design of such a framework and the associated infrastructure involves reconciling objectives which are not necessarily commensurate. That is, the risk framework should:

- Facilitate the efficient and accurate communication of risk perceptions between different audiences in the organisation;
- While providing a consistent and coherent basis for the capture of risk information;
- But without imposing an unnecessary bureaucratic burden on the organisation;
- Or creating a situation where the objectivity (accuracy, fidelity, reliability, etc.) of that information is misleadingly represented.

The CRM's experience suggests that neither "one-size-fits-all" nor "tailor-to-context" approaches can reasonably satisfy all of these objectives at once. Indeed, in light of the discussions in Chapters 5 and 6 it seems likely that these objectives cannot be satisfied through the design of ERM infrastructure alone. Rather, CROs should consider how those objectives might be fulfilled through a range of structural and non-structural interventions.



# Chapter 5

## Rethinking the CRM's approach to decision support

CRM: Risk management isn't something that you come in and set up completely from scratch. Most businesses are already doing things. Good practice, good management is already risk management.. (p 327)

...what we're saying is that there are already elements of this process that we already do. Well of course, there must be. People are making decisions all the time about what's the best course of action for the organisation... (p 328)

*The CRM gives voice to a fundamental realisation (Dialogue 13, pg nos above)*

Chapter 3 described and analysed what Watercare's Corporate Risk Manager envisioned with respect to developing the company's risk management capabilities. Chapter 4 described and analysed the CRM's experiences with implementing that vision. The notable feature was the degree to which the CRM was relatively unsuccessful in his endeavours with respect to his overarching goal of improving the quality of Watercare's "risk data". This chapter describes how the CRM reconceptualised his approach through encounters with two persistent problems:

- **The value problem:** What is the value of detailed "risk data" (and hence detailed risk assessments) in contexts of relatively low uncertainty about system performance or decision outcomes? (Dialogues 11 and 12)
- **The "Holy Grail" problem:** How can CROs translate the abstract definition of the

Risk Management process depicted in world-level ERM standards into a value-adding practice in organisations? (Dialogues 13 and 14)

In each case the CRM proposed a fundamental reconceptualisation of Risk Management and ended up redescribing his role in terms which I characterise as a methodological paradigm shift. The discussion section calls the CRM's "strategic controller" approach (Mikes 2010) into question by framing up the central dilemma of the CRM's experience. Picking up threads from Chapter 4, the discussion then examines the CRM's reconceptualisation by recasting it in terms of the inherent fuzziness of a practical notion of risk. That reframing reveals the CRM's dilemma as a particular instance of a more general dilemma facing CROs everywhere. The CRM's methodological shift is positioned as a proposition that CROs might overcome that dilemma and productively support organisational decision making by adopting a "strategic advisor" strategy (Mikes 2010).

## **Rethinking the role of "risk data" (Dialogues 11 & 12)**

Dialogue 11 opens with the Corporate Risk Manager describing a result from his survey of staff perceptions of the company's risk management function which concerned him (the results of that survey were summarised in Chapter 3). The particular finding of concern was that over 60% of staff did not appear to use the corporate risk register. The CRM interpreted this result as implying that "risk data" was not being broadly used within the business, which, in his own words, "scared the willies out of" him because it implied that potentially 60% of staff believed that there was no need to use "risk data" in their roles. That statistic cut to the very heart of his role as Corporate Risk Manager, which gave the CRM cause to reflect on several questions: "Why are we doing risk management?", "When do we use risk data?", and "What is the purpose of risk data?" In other words, when were risk analyses and "risk data" being used within the business? In this regard, the CRM perceived the use of "risk data" to be somewhat lopsided:

So the question is when do we use risk data? At bottom, what is the purpose of risk data? The only current needs that I know about are the formulation of the Asset Management Plan and capital expenditure applications... (Dialogue 11, p 313)

...it seems to me that there's an easy case for risk data here [pointing to Business Projects], risk data is used to justify business projects. But there's not really a demand or a use for risk data here [pointing to Business Operations]. Risk data isn't used in exercising your current capabilities. (Dialogue 11, p 314)

He explained the perception as follows:

The risk assessment is asking 'what is this organisation trying to achieve with its current operations, what could go wrong?' So, for example, we've got to deliver water, what could go wrong to stop us delivering water? So we do those risk assessments and then if the risk is too large, if it is unacceptable, then that triggers a project. And it's the risk data that is used to justify that the risk is unacceptable. But the point is that we are not using risk data here [Business Operations]. It seems like Operations is where we collect all the underlying data, and we use that data to make an assessment, but that seems to be the only place we're employing it, to justify change, to justify a project. We don't actually seem to be using that data in Operations. The data gets extracted into Planning and then used to justify projects. (Dialogue 11, p 315)

In these passages the CRM revealed a perception that "risk data" was apparently not widely employed in the execution of what he called the company's "current capabilities". That is, Watercare's risk management function did not appear to provide much support for operational decision making, in contrast to the explicit and widespread use of risk assessments and "risk data" to support capital decisions (i.e. to justify projects). This was not entirely accurate, since, as I pointed out in Dialogue 11, formal risk assessments were visible in certain operational contexts. For instance, "risk data" was employed in the form of reliability modelling to justify the maintenance programme (although, strictly speaking, that was a planning function), basic risk registers were used to inform the management of projects, and the company had a rigorous health and safety programme which involved the formal assessment of hazards. Nevertheless, the CRM had a point. While risk was explicitly a key factor in *all* capital decisions, it seemed that, with the exception of few isolated contexts, risk was not employed to the same extent to support decision making by the Operations group (see Figure I.3, Appendix I, page 352).

This observation was somewhat perplexing for the CRM since the risk assessment process was ostensibly a means of analysing current capabilities to establish the need for change, and, at least in principle, could be applied to any process in the organisation:

I mean, Operations is about your current capabilities... Projects are about changing your capabilities... And, in order to work out what capability you need for the future, you need to understand what your current status is. At least in Watercare, projects, I mean the need for a new capability is driven by a realisation that your current capability is going to be short at some time in the future, or its not going to be adequate, or something like that. So somehow you must be analysing what you're currently doing to recognise that need. (Dialogue 11, p 314-315)

This was precisely how "risk data" had been used to support decisions about plant and

equipment maintenance within the Reliability Centred Maintenance programme:

..What are they doing? In essence they're taking risk data and they're using that risk data to influence their decision making about maintenance. So they've identified risks that could result in kit breaking down, they've got information on how the kit could break down and the relative importance of that, and then they decide how to manage that kit so that it doesn't break down. Is that a form of treatment, or management, as opposed to change? So in other words, is that decision making as opposed to change? Now I don't know that the two things are different, do you know what I mean? (Dialogue 11, p 315)

I subsequently commented that there appeared to be a correlation between the magnitude or significance of decisions and the use of formal risk assessment within the organisation. That is formal risk assessments were clearly used to support relatively infrequent but significant decisions about capital expenditure (and in the case above to manage a large programme of operational expenditure), but were not used to support every-day decision making by individual actors in the company. This latter sort of decision-making by actors in the performance of their roles constituted the majority of decisions made within the organisation on a day-to-day basis, most of which would have been made many times before under the same or similar conditions. Such decisions were thus routine to the performance of existing capabilities, as the CRM put it, and were not generally made under conditions of significant uncertainty about outcomes.

It was in this regard that the CRM felt the finding from his survey, that “greater than fifty percent of respondents have no formal knowledge or awareness of the risks that could impact the performance of their business unit”, was significant (Dialogue 11, p 316). The CRM commented that the result was “a concern because it may indicate that the staff follow business processes because they're supposed to rather than because they understand what the business processes are designed to achieve in terms of risk control” (Dialogue 11, p 316). The CRM's point was that, in many cases, existing business systems and processes were risk controls; for instance, health and safety procedures, induction processes, and procedures for managing financial transactions were all extant mechanisms to control risk. But while such procedures and processes may originally have been conceived and designed as risk controls, over time and for various reasons the original purposes may be forgotten such that they simply become part of the way things are done. The CRM's concern was that without an explicit understanding of why those procedures exist, actors might be inclined to cut corners, perhaps giving only token consideration to formal procedure, or even ignoring it altogether.

The CRM's concerns may not have been entirely justified (i.e. it seems unlikely that the majority of Watercare's staff blindly followed procedure without understanding what they

were doing), but this line of on-the-fly analysis pointed up an important function of formal risk analysis as a means to achieve “...insight into what we already do... you could argue that it gives you a better understanding of why you do things” (Dialogue 11, p 316). In other words, the value of performing formal risk assessments in the Operations context was not so much as a means to identify risks to the performance of existing capabilities, since these would mostly be negligible, but rather as a procedure for reflective practice (Schön 1983). As the CRM put it:


...there's value in that, in that it will hopefully encourage people to think about what they're doing. And if they think about it then perhaps they can change it for the better. So that's what we're saying, by making decisions, any decisions, based on risk data, we will make better decisions. If we compare what we're doing against risk data it should give us a basis for saying whether or not we could do things better. (Dialogue 11, p 316)

In Dialogue 12 the CRM returned to the question of the value of formal risk assessment. He commented that the process of designing the risk framework was, in essence, one of making explicit what the organisation already did, of drawing “a picture using mathematics” that “captures and conveys the reality” of “work that's already done in the business” (Dialogue 12, p 321). The CRM felt that there was a need for this, since the existing framework was a poor reflection of Watercare's existing corporate objectives and performance standards (“it's not representing very well what the business already does”; Dialogue 12, p 319), and that there would be value, in terms of the defensibility of decision making, in achieving a more objective representation of reality:

...presumably there is value in the sense that that is precisely the role of risk management from a governance perspective, to make sure that the Board is getting a realistic representation of the company through the risk analysis and data... So I think the practical need there, and the value that I see my work adding to the business is with helping to formulate the AMP. If I can get the risk data to a point where they can analytically calculate the AMP scenarios then presumably the AMP will be a lot more defensible than at present where the decisions on which projects to cut are somewhat more subjective. So at a higher level I tend to think there's some kind of defensibility value there as well, which you could paint as providing a more objective basis for decision making. (Dialogue 12, p 321)

But the CRM was less confident that he would be able to provide value in terms of “providing insight into what the business doesn't do well” because Watercare was a mature organisation:

...part of me wonders about the fact that I would call Watercare a mature



organisation, in the sense that the company, in whatever form, has been around for the better part of half a century. If you look at it from an organisational knowledge perspective, where the business develops and retains knowledge about how it goes about its business, then presumably, over time, the organisation has developed better and better ways of doing things. Now Watercare's been around a long time, so when I look at the business and what it does I believe, and I have believed right from the beginning, that they're pretty good at managing the common risks associated with the business. Sure, there are probably little holes in amongst what they do, where the company might be exposed, but in terms of the core stuff I think they're probably pretty good at it... (Dialogue 12, p 319)

In this Dialogue, the CRM's use of the term "framework" was somewhat broader than its use in Chapters 3 and 4. Those chapters were concerned specifically with the physical document which specified the categories and evaluation criteria by which risk was to be assessed within the company (at Watercare this took the form of a table which defined evaluation criteria across various consequence categories and levels of severity). In Dialogue 12, however, the CRM was referring more broadly to the idea of making explicit, in a representational sense, the connections between the company's existing business systems and processes and the various objectives and risks they were respectively intended to achieve or control. Such a framework would, in effect, constitute a hierarchical risk register (as the CRM had envisaged in Chapter 3), the definition of which would amount to performing a risk assessment on the enterprise as a whole. As the CRM conceived it, the purpose of such a framework would be, first, to accurately reflect "what the business currently does", and then, ideally, provide "insight into what the business doesn't do well". The CRM's concern was that because Watercare was a mature organisation, the category of what the company "didn't do well" was probably very small, which called into question the value of such a detailed deconstruction of the business systems and processes:

...Now you could say the purpose of the risk management function is to analyse where there are holes, but this is a mature business so what's the chance that I'm going to find big holes? In which case, what's the value? Is it worth spending my salary every year plus the distraction I cause to other staff just to find what are likely to be little holes in the organisation? (Dialogue 12, p 320)

For the CRM, the challenge presented by Watercare as a mature organisation<sup>13</sup> was that in order for the risk framework to provide "insight into what the business doesn't do well",

---

<sup>13</sup> The CRM's comment that Watercare was, on the whole, "pretty good at managing the common risks associated with the business" appears contradictory with respect to his conclusions about the company's Risk Management function and his assumed decision support rationale (see Chapter 3), and indeed with the results of the staff survey discussed in Dialogue 11. This apparent contradiction is examined in Appendix VII.



the framework would have to be defined to a very high level of detail:

I think that Watercare is a fairly mature organisation and I think you're going to have to work at a very high level of detail to provide additional insight into what they do. So I think it's more challenging in a mature, well developed, well experienced organisation, for the risk management function to provide additional insight because it involves having a very complex or very detailed look at things in that kind of organisation. (Dialogue 12, p 322)

Watercare's Reliability Centred Maintenance programme stood as an example of the CRM's point. As a mature organisation, the company already had well established maintenance routines, but the RCM programme has been promoted and justified on the basis that reliability modelling would identify inefficiencies in those routines, at least for Watercare's above ground assets (i.e. treatment plants and pump stations). In order to do this, however, those models had to be defined to very high levels of detail, right down to the smallest discrete components of individual pieces of plant and equipment. The CRM's concern was that such detailed analyses were too costly and difficult to perform anywhere else.

The resulting dilemma was that which was already described in Chapter 4. That is, taken to its ultimate extreme, the CRM's goal of an objective risk framework would amount to representing all of the organisation's objectives, performance standards and parameters in absolute detail, and in a format that made explicit all of the relative priorities between those objectives and standards. It would also involve making explicit all of the organisation's existing risks and risk controls (i.e. how those risks were controlled by existing business systems and processes). In essence the framework would mirror the organisation. But such a framework would, of course, be utterly impractical both from a design perspective and in use. The other end of the spectrum was, however, equally undesirable, where the correspondence between the risk framework and the company's objectives, performance standards, and processes would be so loose as to be practically useless for internal audit or control. Conceptually, a suitable risk framework must lie somewhere between these two extremes, representing a balance between objectivity and practicality. The practical question was, therefore, how detailed should the framework be?

As in Dialogue 11, the CRM subsequently expressed the opinion that the potential value of formal risk assessment might therefore lie not so much in identifying "holes" in organisational practices, however "big" or "little", but rather in providing an opportunity for actors to reflect about the way things are done:

...what you might find is that by forcing the engineers to sit down and think about

the various sources of risk and to think about the actual business systems that are in place to mitigate different areas of risk, what you might make them realise is why they do things... Part of me wonders that if you sit down and analyse the business and why it has these practices what you might realise is why... (Dialogue 12, p 322)

...so if you're interpreting a business practice in terms of a risk management system... you might better understand the purpose of that business process. And if you understand the purpose of that business process, then you might appreciate the relative importance of different business processes, and if you understand the relative importance of different business processes then you might change the way you behave, I guess. As opposed to sending a graduate into a room and saying "do those calculations, send out those documents, draw those drawings", and he does them and when he's short for time, he cuts corners wherever he can. If, after a year of doing that kind of work, you said to that graduate "look at what you do and why you do it in terms of risk management", they might say "actually, producing the drawings at the end is the most important thing, so if I'm going to cut corners I'll cut them at another point", I guess. So, I'm not sure, but in terms of a mature organisation you're not necessarily providing new insight. In some ways all you're doing is providing justification or almost prioritisation of what they're already doing. But the bottom line in some ways is that concept that you've mentioned to me before, that the risk analysis process is almost making them reflect on what they do and why they do it... (Dialogue 12, p 323)

## Rethinking Risk Management (Dialogue 13)

Dialogue 13 was a conversation that I had with the CRM about a presentation that he was preparing to deliver at an ERM conference. The title of the presentation was "In search of the holy grail: practical challenges in the pursuit of ERM". In the first part of Dialogue 13 we threw around a number of ideas, in the form of challenges encountered, lessons learned, and observations made over the previous 12 months. These included:

- Key ERM concepts expressed in world-level standards were: "integration, business processes being embedded, being enterprise wide, full depth, full breadth of the organisation, holistic, all sources of risk and continuous, forward looking". (p 325)
- That Watercare's stand-alone risk register was an example of "risk management being detached from the day-to-day business". It was effectively "bolted-on" as an extra layer of bureaucracy rather than being integral part of business processes. (p 325)
- That the risk register is, almost by definition, a "bolt-on" solution because the register merely captures statements of risk which are the end-products of risk assessment processes. The register doesn't support the production of knowledge about risk but

“simply sits as a recording device at the end of that process”. (p 326)

- That there are many ways to measure business performance, and the process that the CRM went through to redefine Watercare’s risk assessment framework was really a process of working out what mattered in terms of performance: “...you’ve got a raft of possible ways you could measure organisational performance. You end up going through a process where you’re trying to boil it down to a handful of those that matter”. (p 325)
- That the concept of risk management being enterprise-wide puts a focus on how the organisation as a whole performs and therefore “focuses everything on a key driver, a common driver, which requires you to think about what the principal objectives are for the organisation”. (p 326)
- That the concept of risk management being full-depth means “you need risk management activity at lower levels to be correctly aligned and supporting decision making at the top”, and this “marrying top down with bottom up” requires “aligning responsibilities for risk with those who can actually control it” and an efficient “translation of information from the bottom of the organisation to the top” (the risk framework was supposed to act as a mechanism for this translation). (p 326)

The preparation of the conference presentation was the really the first opportunity the CRM had had to step back from his day-to-day work and reflect on all of his experiences from the past year. But, the CRM wanted to do more than simply describe a few challenges faced or lessons learned in implementing ERM at Watercare. As he put it:

...if I just get up give three examples of challenges that we’ve faced, and then walk away, I haven’t really added much. Sure I can make the point that it’s complicated, it’s hard, but I’ve not really added much, I haven’t really pulled anything together. (Dialogue 13, p 327)

Rather, the CRM wanted to address what he referred to as the “method gap” in the international literature with respect to how to turn the generic concepts from ERM standards into practical solutions within the business:

Certainly, when I started to try to explain these problems, or even just to put these issues into words, I realised that I was being drawn back to the fundamental question of what are you doing here? Fundamentally, what are you trying to achieve? What do you need data for? What is it that you’re supposed to do? What role does the data play? Because the risk management process is basically the process for decision-making, it closely resembles the rational decision making model. But Enterprise Risk Management is also talked about as a

culture. So this is the question: how do you go from what is essentially a business process definition to a culture? I don't know. (Dialogue 13, pp 327–329)

Ostensibly, the key methodological question was how to integrate Risk Management (represented by the abstract process definition in ERM standards) into the business such that it would become part of what organisational actors did everyday; “so that it underpins what they're doing”, as the CRM put it (Dialogue 13, p 328). This principle of integration is a distinguishing feature of ERM, referring both to the cultural embedding of Risk Management throughout the organisation and its processes, and, in a calculative sense, to taking a cross-functional, enterprise-wide perspective of risk (Lam 2003, p 45; see Appendix V). Enterprise Risk Management standards and capability maturity models emphasise the normative importance of this integration. The advanced ERM capability, for instance, is normatively defined by criteria such as “institutionalised” and “permeating culture”, which are intended to indicate the evolution of Risk Management from merely a formal process, “bolted-on”, so to speak, to a culturally integrated practice underpinning organisational decision making. But, in the CRM's opinion, the world-level literature provided little or no practical guidance for achieving this transition, which was a method gap that he wanted to address by “pulling together” his experiences at Watercare (Clement 2008b, Slide 10).

Reflecting on his experiences over the previous 12 months the CRM emphasised a key understanding that the implementation of ERM was really about interpreting the business from a risk-based perspective, and that this required a common basis for the comparison of business activities:

...really what I've done is to have a look at our existing business practices, and the fact that we have these existing features here, so for example the health and safety management system, the compliance management system, and so on, these are all evidence of what could be interpreted as risk management activities. In essence, I realised that if I was going to be efficient in what I was doing, I shouldn't try to re-engineer the business but to integrate risk into it. It's about interpreting what these things are under a title of risk, if you like, finding a common basis for assessment across the business, finding a way to compare apples with oranges, that's this whole thing about risk being a common basis on which to make a comparison. So some one could come in cold, set up a new system and say “I'm going to apply my system to your business and compare things in your business based on my system”. Or you can come in and look at the business and say what here would give me a good basis for common comparison. And, really, that basis is the corporate objectives, and what risk management tries to do is to compare everything that goes on within the organisation on the basis of how it could impact on the achievement of those objectives. So I think it's about interpreting the existing business practices from the risk management perspective, but coming up with that framework is really part art, part science. You need to be able to look at a business and understand it, and at the same time

you need to be able to look at your objectives and deconstruct them, and then try to marry the two together. You have some paper-based objectives, and then you have real business practices and what you're trying to do is to bring these two together to achieve an efficient translation of information from the bottom to the top. (Dialogue 13, pp 326-327)

At the time Dialogue 13 took place the understandings in the above passage represented a substantial learning curve for the CRM, and were not yet fully crystallised in his mind. As the CRM commented, "Granted, there's no simple solution, which is I guess where this bit about it being as much art as science comes in, but that's pretty vague" (Dialogue 13, p 327). A key understanding that was clear, however, was that much of what the business already did was risk management. It was not just the case that Watercare already had a formal risk framework and register, or that explicit risk assessments were already performed in various contexts within the organisation, but rather, and more fundamentally, that the company's existing systems and processes could be interpreted as risk management activities. This was the same point that the CRM had already made in Dialogue 11 when he observed that Watercare's existing business systems and processes could be seen as risk controls, and in Dialogue 12 when he depicted Watercare as a mature organisation that was "pretty good" at managing the common risks associated with the business.

The above understanding was the result of viewing the organisation through what the CRM would later refer to as a risk management lens (see Dialogue 14). In this regard, the CRM's exercise of marrying together "paper-based objectives" with "real business practices" was, in effect, a form of enterprise-level risk assessment from which perspective certain organisational systems and processes were clearly identifiable as controls intended to constrain or mitigate the potential for undesirable outcomes in the company's internal and external environments (even if they were not formally referred to as "risk controls" within the company). When the organisation was viewed through such a lens it became apparent that risk was managed not just through explicit performances of the Risk Management process, as that process is depicted in international standards, but through many of Watercare's usual business activities:

Risk management isn't something that you come in and set up completely from scratch. Most businesses are already doing things. Good practice, good management is already risk management... (Dialogue 13, p 327)

...what we're saying is that there are already elements of this process that we already do. Well of course, there must be. People are making decisions all the time about what's the best course of action for the organisation... (Dialogue 13, p 328)

Strictly speaking, in our impromptu discussion in Dialogue 13, the CRM and I did not properly distinguish how or to what degree Risk Management at Watercare was already constituted as formal process versus cultural practice. Rather we intuitively (and somewhat uncritically) placed explicit and therefore observable instances of formal risk management activities (e.g. risk assessments, risk reviews, risk reporting etc.) into the category of “Risk Management as formal process”; while the rest of the company’s business activities were subsumed under the category of “Risk Management as cultural practice”. In essence, our distinction was based on the explicit use of the term “risk”. That is, all activities in the former category involved the explicit use of the term “risk” by the actors involved in performing those activities. The actors were aware that they were performing some sort of risk management activity, they talked explicitly about risks, risk controls, and so on, and the activity was carried out under the umbrella of Watercare’s Risk Management Framework. In contrast, the activities in the latter category did not involve the explicit use of the term “risk” by the actors involved, or if they did, its use was incidental to the purpose of the activity. Whether the activity was a formal process or procedure, or simply some ad hoc practice in a certain context was irrelevant to our distinction. Activities in this latter category were, rather, “interpreted as risk management activities” as a result of viewing the organisation through a risk management lens.

An example of such an activity at Watercare was the Shutdown Register. In one meeting Watercare’s project managers described how they had developed a register to track and coordinate planned shutdowns within the water and wastewater networks. Such shutdowns were necessary for maintenance and for connecting new equipment to the existing network. The register had evolved in an ad hoc fashion from a simple tabulated list to a more formally constituted programme, supported by a documented procedure for communicating with the Operations group. The register had evolved into a well-developed method for controlling an important source of uncertainty in project management work, but it was not the product of any formal Risk Management process. Even though it had since been formally recognised as a risk control in Watercare’s corporate risk register, the development of the register was an example of what we called “Risk Management as cultural practice”.

For the CRM, the “lens” metaphor reflected a vision of risk management being so ingrained within the organisation culture that it is no longer “bolted on” to the business as an imposed administrative task, but rather naturally performed as a way of doing business:

Finally, ERM means that the risk management process has been effectively and totally integrated in the business, so that the risk management activity does not adversely impact business efficiency. Risk management simply becomes ingrained in the day to day process of doing

business – its ceases to be an additional thing we have do, and instead becomes a way of doing business. Practically this means that the risk management process ceases to be analytical and instead becomes intuitive because the organisation's staff have all the necessary information, have all the necessary competencies, and clearly understand the organisations current appetite for risk. (excerpt from Clement 2008b, presentation to ERM Conference, April 2008, Slide 8)

In his presentation, the CRM analogously compared the achievement of such an “ERM culture” to attaining the Buddhist state of Nirvana, this being the “holy grail” referred to in the title of the CRM's presentation:

In Buddhism nirvana is a state of peace and enlightenment that involves being unaware of one's self. It is the liberation from suffering and departure from the endless cycle process of reincarnation.

In business Enterprise Risk Management (ERM) is the state of performance reliability and efficiency that involves the organisation being unconsciously competent. It is the liberation from risk administration and the endless cycles of formal process-driven risk management.

(excerpt from Clement 2008b, presentation to ERM Conference, April 2008, Slide 9)

The CRM explained this in Dialogue 13:

...I'm presenting material that only barely scratches the surface of some of the issues that I've been grappling with of effectively achieving the “nirvana” state of ERM. Now I never knew really what I was thinking there, but when I looked up what “nirvana” means, in Buddhism nirvana is a state of peace and enlightenment that involves being unaware of one's self, and the phrase that was used was being “unconsciously competent”. Another definition was that it's the Buddhist term for self-realisation, the transcendence of suffering. Well, going through the risk management process is a process of suffering, it's an administrative process, and what we hope to get to is the transcendence of that suffering where you have complete self-realisation and you're unconsciously competent. (pp 328–329)

For the CRM, the “nirvana” analogy was a way of interpreting and making sense of the challenge of integrating Risk Management within the organisation:

To me this analogy summarises the vision of ERM – the organisation is efficient and has reliable performance because it is essentially unconsciously competent in what it does. It doesn't require an imposed analytical framework to make sure that the right decisions are made for the organization. Rather, good quality decision making practices are essentially ingrained into the culture and knowledge of the organisation and the selfless behaviour of people within it. (excerpt from Clement 2008b, presentation to ERM Conference, April 2008, Slide 9)

Whether the “nirvana” analogy was a suitable device for making sense of the concept of “ERM as culture” is debateable. The term “unconscious competence” is suggestive of Ryle's

notion of skillful or intelligent action (1949), but it is also strangely contradictory in the sense that neither a competent person nor a competent organisation can be said to perform “unconsciously”. A degree of conscious awareness is required for intelligent action (see Chapter 6), and, indeed, rigorous, formal process is often essential to competent decision making, whether personal or organisational. Furthermore, despite the attractiveness of the idea of being liberated from bureaucracy, administrative procedures will always have their place and purpose in organisations. On the other hand, however, the “nirvana” analogy did have parallels with cyclical concepts of learning (such as Nonaka and Takeuchi’s SECI model, already described in Chapter 4; Nonaka and Konno 1998; Nonaka and Takeuchi 1995). I commented on this in Dialogue 13:

Hmmm, that’s interesting, there are multiple parallels there. We’ve got “culture vs process”, that’s one way of putting it. Another phrase we used was “implicit vs explicit”. And still another way of looking at it, taking an analogy from the knowledge management literature, is “tacit knowledge versus explicit knowledge”. So, what you just talked about there, the process of suffering leading to the nirvana state of self-realisation, that seems to describe the process of organisational learning. So you tacitly understand what you’re doing, it’s intuitive. Then you go through an explicit process of documenting that knowledge and reflecting on it, so that’s the risk analysis process, the process of suffering, which then leads you to further tacit knowledge. It’s a cycle. (Dialogue 13, p 329)

These parallels point up that “Risk Management as formal process” and “Risk Management as culture” are not so much two developmental states of ERM but, rather, should be seen as two complementary characteristics of a mature ERM capability (it is in this sense that the CRM’s “nirvana” analogy was perhaps a reasonable caricature of the ERM vision). The generic risk assessment process, as it is depicted in world-level standards and frameworks, is explicitly a procedure for looking at the organisation and re-interpreting existing business processes and practices from a risk management perspective (viewing the organisation through a risk management lens). It implies a methodology for investigating the organisation and determining how the various business systems and processes contribute to the achievement of the corporate objectives and the control of threats to performance. Through such inquiry new and deeper understandings – new knowledge – of business activities may be acquired and passed on; which, according to the ERM rationale, is assumed to be the wellspring of organisational performance. But such inquiries must, of necessity, be explicitly performed. Even if the methods employed lack formal rigour, one cannot conduct an inquiry without explicitly and therefore consciously intending to do so, and knowledge cannot be transferred between actors without being made explicit in some form, whether



representational or instructional (Leonard and Swap 2004; Nonaka and Konno 1998; Nonaka and Takeuchi 1995).

Thus, the analytical aspects of the risk assessment process cannot be escaped, but, to the contrary, are essential to the production of knowledge. In this sense, “Risk Management as culture” does not refer to some transcendental state of enlightenment in which the organisation is somehow liberated from the need for “Risk Management as formal process”. Rather, it refers to the idealised aspiration of ERM that the tendency to see things through the “risk management lens”, and hence the tendency to explicitly assess and account for risk in decision making, should become ingrained or culturally embedded within the organisation. It is not the escape from “suffering” but rather the realisation that one cannot escape it: formal procedures and methodologies for the assessment and management of risk cannot achieve much if organisational actors are unaware that they should put them into practice; and, vice versa, a cultural awareness of risk cannot achieve much without the appropriate procedures and methodologies for its assessment and management. Thus, “Risk Management as formal process” and “Risk Management as culture” must be seen as essential and complementary aspects of the mature ERM organisation.

## The CRM rethinks his approach (Dialogue 14)

Dialogue 14 was the last formal interaction between myself and Watercare’s Corporate Risk Manager during the main data collection phase of the research. At that time the CRM had been with the company for approximately 15 months so I deliberately asked him to reflect back over that time and to describe how his approach to his role had evolved. The CRM described how he had come into the role with an initial expectation, rooted in his background and training as an engineer, that there would be some “purity in the analytical approach” (Dialogue 14, p 332), that risk could be measured and quantified using robust analytical methods. Over time, however, he had come to realise that there was “no perfect analytical method” (p 332), that all methods are subjective to some extent, and that all relied on the adopted specification of risk, which itself would always be just one of a multitude of possibilities for how risk could be defined and measured. The CRM’s initial approach and his subsequent experiences, which caused him to doubt that approach, were analysed in Chapters 3 and 4.

Those experiences led to two key realisations. The first was that the business already managed risk (“the business already does risk management”), and thus Risk Management

was not “some kind of new process that the organisation is going to start doing”, but rather was “simply a way of looking at the organisation, at what the business already does” (Dialogue 14, p 335). This was the same reconceptualisation of Risk Management that the CRM had tentatively suggested in Dialogues 11, 12, and 13, but here he expressed it much more confidently and concretely by explicitly referring to Risk Management as a “polarising lens” (p 335) for interpreting the business in terms of threats and controls. The second key realisation was that Risk Management “may not necessarily fulfil its roles by being analytical” (Dialogue 14, p 332). While analysis would always be implied and necessary to an extent, it did not necessarily “have to be quantifiable” (p 332).

Together these two realisations constituted a fundamental shift in perspective for the CRM and formed the basis for his reconceptualisation of his role. He revealed how his thinking about the role had consequently undergone what could literally be described as a paradigm shift. As regards the purpose of the role, he had evolved from a narrow focus on generating good quality (objective) “risk data” to a broader and perhaps less well defined understanding that his role was primarily to:

...influence thinking, to influence the way that people think about the organisation, the way that they see the organisation internally, how they see their internal customers, and how they might view their relationships with other people in the business, so that they have better clarity about what it is they're required to do and why they are required to do that. So, at this point in time, I think my role is primarily concerned with influencing how people understand what they're doing and why they're doing it, which is the cultural change that I'm trying to achieve. (Dialogue 14, p 331)

This shift in focus was accompanied by an important shift in methodology. The CRM felt that the reconceptualised role did not require a lot of precise analysis and detailed data (which anyway could not always be achieved even within the technical contexts of the Watercare business), but rather required someone with insight, someone “who can look at systems, patterns, interactions, and who can cut to the chase” (p 333), and who could then productively engage organisational actors in reflective dialogue. As far as the CRM was concerned, this was where the core value of his role lay:

...I don't get positive feedback from the business about the data, like “Oh yeah that data's really good”. At the end of a meeting, people say “that was really interesting”, or “that was a really good meeting”. But it's only because people have been around a table, talking and debating, and it's changed the way they've been thinking, or they've made some progress or got agreement on something. So I think that the organisation really does get value out of reflecting on what it does, and out of having someone like me critiquing the organisation and talking to people about it, and trying to get them to think about what they do and how

they do it, or getting them to view the organisation in a different way to how they normally see it in their day to day business, to stand back and look at the big picture. And you don't need data to do that. (Dialogue 14, p 333)

The CRM now believed that he could deliver more value by directly engaging with actors in order to influence perspectives and behaviour, and that he could perform this role largely without the detailed analyses and “risk data” he had so fervently sought after. As he put it, “So my role is to influence thinking, and influencing thinking may not actually require data generation” (Dialogue 14, p 332). Despite this 180° turn-around in his approach to the role, however, the CRM continued to perceive that data generation was an expected function of Risk Management:

...I think the data is expected. So people, both inside and outside the company, expect to see risk data if you're doing risk management. They expect to see a risk register with measures of things... I would expect that if [the external auditor] came in here next year and said “show me your risk information”, and all I had was a list of issues with no explicit measurement framework, I think they would criticise the method because it's not analytical. (Dialogue 14, p 332)

In addition, there remained a tension with the governance and internal control aspects of the role. That is, CRO's still have to report to the Board and require an objective basis on which to do so. As the CRM put it, “How do you provide the Board with an indication that things are well managed if someone is just floating around, trying to facilitate development where they think it's needed? I guess that's why you need to have data in the end” (Dialogue 14, p 334). Also, CROs still need to defend their role, and can't do this without data: “The irony there is, I'm starting to see that the data, the process of generating a lot of detailed data, is probably, not a waste of time, but a huge inefficiency, but then, without it, how do you justify that you're focussing on the right issues?” (Dialogue 14, p 334). In this regard, the CRM had come to understand that his role required someone who was both an analytic and synthetic thinker, “...there's a bit of an art to it, it's not pure science, it's not well defined black and white” (Dialogue 14, p 336). Risk Management implies the analysis of risk but, as the CRM pointed out, the critical question is “...how much detail is appropriate?” (Dialogue 14, p 335). Risks need to be analysed, and described, and where appropriate quantified, but that analysis has to “marry up with the specific needs of the business” (Dialogue 14, p 335). The CRM referred to this “marrying-up” as the art side of Risk Management:

...so it's made me realise that's the art side of risk management. I mean, my framework is explicit knowledge, an explicit representation of the business. And sure, I've gone through a process of analysis to generate it, but it's relatively crude and I don't believe it is as influential, in terms of influencing

behaviour, as looking at the organisation, recognising where there are problems, seeing how things interact, and then going to someone and saying "did you think about it this way?"... So the art is very much, I think, about how you marry the analysis with the way the organisation thinks. (Dialogue 14, p 336)

## Discussion: "strategic controller" in question

Chapter 3 described how Watercare's Corporate Risk Manager interpreted the task of implementing ERM as essentially one of "design and build". That is, the CRM's approach reflected an interpretation of ERM as essentially a new capability that needed to be integrated into the organisation by literally designing and building the necessary infrastructure (i.e. risk frameworks and registers, and calculative tools), and training organisational actors to perform "Risk Management" (i.e. creating the so-called "risk culture"). The assumptions were that "good" decision making required the formal assessment of Risk (and therefore formal Risk Assessments should be performed for all decisions), and that the primary goal of those assessments should be the production of what the CRM referred to as good "risk data" (i.e. objective, quantitative representations of Risk).

Chapter 4 problematised the CRM's "design and build" approach to implementing ERM by revealing the dilemmas associated with defining detailed risk frameworks and registers. It was argued that what the CRM did, *in effect*, was to impose on other actors his own perceptions about how much they needed to know (about risks) in order to decide and act in a given situation. Chapter 4 also called attention to the cost, in terms of the time, resources, and calculative expertise, of producing good "risk data", and revealed the difficulties which arose out of the conflict between the CRM's imposed standard of data quality and what was realistically possible in practice. This raised the perplexing question of value; i.e. at some point the cost of pursuing better "risk data" might outweigh the potential benefits of more informed decision making.

Dialogues 11 and 12 framed this cost-benefit dilemma with respect to what the CRM referred to as the "performance of current capabilities"; both at the level of individual actors and their day-to-day practices, and at the level of the enterprise as a whole. The specific issue was that it was not clear that formal, detailed risk assessments would actually add value in contexts of relatively low uncertainty about system performance or decision outcomes. For instance, in performing their roles day-to-day most actors are probably pretty competent, and confident, and the risks associated with those individual performances will be negligible. The same basic logic applied to the CRM's consideration of Watercare as a mature

organisation. Even where the consequences of mistakes or failures are high, and precisely because they are high, the likelihood of such consequences actually being realised would be low due to existing risk controls. The CRM had realised that, in order to fulfil his idealised standard of objectivity, any formal assessment of risk in such contexts would have to be performed to a relatively high level of detail to distinguish and quantify risks, but also that those risks would not likely be significant (what the CRM referred to as “small holes”). In other words, the costs of performing detailed risk assessments in contexts of relatively low uncertainty were likely to outweigh any benefits that the resulting “risk data” might have for decision makers. Thus, two difficult questions emerged to constitute one of the central problematics of the CRM’s experience:

- 1) If objectivity is costly, then what is the value of seeking better “risk data”, and how can that value be demonstrated?
- 2) In some cases it may not be possible or cost-effective to obtain “good” data. What should happen in these cases? How can Risk Management fulfil its decision support and assurance functions in such situations?

In Dialogue 12 the Corporate Risk Manager expressed a view of Watercare as a mature organisation which consistently performed well in its local context, and which arguably was “pretty good at managing the common risks associated with the business”. The CRM’s perception was supported by consideration of the firm’s institutional and operating contexts, and long-term enterprise performance (see the overview in Appendix I). However, this view appeared contradictory to the rationale which underpinned the CRM’s original concerns about Watercare’s “risk data” and his subsequent development proposals (see Chapter 3). That is, it was the CRM’s assumed rationale that “good decisions” should be based on “good data” which rendered the quality of Watercare’s “risk data” problematic in the first place. Analysis of the apparent contradiction (see Appendix VII) concluded that Watercare’s “risk data” was in fact consistent with how actors within the company actually understood or perceived risk. Indeed, if it were not then Watercare’s managers and staff would not have explicitly referred to the register as the repository of information on the company’s identified risks, nor referred to graphical representations in reports and Capex requests as the primary means of communicating the significance of those risks. But while that “risk data” may have been sub-standard, as judged against certain abstract expectations of what risk should look like, it was apparently not problematic as far as organisational decision making was concerned. Watercare’s managers and employees may not have been able to specify and quantify risk in terms consistent with the theoretical definition (i.e. as a quantified

probability distribution), but this did not stop them making good or at least adequate decisions, as indicated by the long-term performance of the enterprise. In other words, what Watercare's management and staff knew about the company's risks was appropriate to their respective decision making purposes, at least most of the time. If this is accepted then the root of the above contradiction lay not with the rationale that more informed decision makers generally make better decisions, but rather with the implicit assumption that the reverse is also true, i.e. that "good" decision making requires a certain quality of information or knowledge. The discussions in Dialogues 11 and 12, and the analysis in Appendix VII, thus call attention to the fact that high levels of precision and quantification may in fact be unnecessary with respect to the pragmatic requirements of deciding and acting in certain contexts. The CRM's efforts to justify his proposals for improving Watercare's "risk data" were therefore rendered further problematic by a third question:

- 3) How much detail is necessary with respect to making good decisions in various contexts, and providing stakeholders with assurance that the organisation is effectively managing uncertainty? How can the appropriate level of detail be determined?

At the time I think the CRM and I both perceived these questions to a certain extent, even if we could not make them explicit. Rather, they lingered in the background of our endeavours as a vague uncertainty about the burdens of pursuing "good data". That we did not immediately see them as directly challenging our rationales about the implementation of ERM was probably due to an intuitive reluctance to contemplate the idea that the pursuit of objectivity might be the wrong approach. In hindsight, the CRM was never in a position to answer these questions, due, fundamentally, to the fact that the answers depend on the particular relationship between the decision maker and the contextual domain in which the decisions are being made (this is explained in Chapter 6). It was this fact which ultimately prevented the CRM from realising his vision of Risk Management as a specific capability to be developed within the organisation, i.e. as a formal procedure to be followed when making decisions, and requiring the use of specific tools and calculative infrastructure.

The CRM eventually reconceptualised his approach in order to overcome the problems that he encountered. As was explained above, that reconceptualisation was grounded in two key realisations: (i) that the business already managed risk ("good management is already risk management..."), and thus that Risk Management was not a new process that needed to be "bolted on" or even integrated into the business, but rather was "simply a way of looking at the organisation, at what the business already does"; and (ii) that while some form and degree

of risk assessment would always be implied and necessary, it did not necessarily “have to be quantifiable” in order for Risk Management to fulfil its roles. The following section examines why both of these realisations were necessary in order for the CRM to overcome the dilemma posed by the above questions.

## **Reinterpreting Risk Management**

The analysis in Chapter 4 drew on a meta-theoretical discussion of the ontological and epistemological status of risk. First, risk was defined, in concept, as referring to the possibility of future loss (or gain) where the expected outcome is only more or less likely but not certain to occur over some future time period. Second, risk was defined as having both real and subjective components. That is, as a label risk does refer to real states of the world, but, at the same time, what counts as significant, and therefore as risky, is the product of social value judgements – some state of the world must be recognised and judged by someone in order for risk to actually exist. Thus, any statement of risk expresses expectations about three things: (i) a future state of the world resulting in certain outcomes, (ii) how significant (good or bad) those outcomes will be, and (iii) how likely or unlikely those outcomes may be within a specified period of time. Third, drawing on Knight’s (1921) original distinction, risk was distinguished as a specific form of knowledge about the future where there is some objective basis for quantifying the probability of expected outcomes. That is, in order to specify risk we must know enough to reasonably calculate and specify the real probabilities of outcomes; where we cannot make that estimate then we are merely uncertain.

But while such a discrete boundary between quantitative/qualitative knowledge is easily conceptualised in the abstract, it becomes much fuzzier in practice. Any estimate of the real probability of outcomes, technically referred to as aleatory uncertainty, is to some extent uncertain in an epistemic sense because our knowledge of the future can never be perfect. This epistemic uncertainty can never be eliminated, but merely reduced through further inquiry to the point where it may be safely ignored for the purposes of calculation. In many decision making contexts epistemic uncertainty may be trivial and therefore practically irrelevant, but definitionally the implications are significant. If what we know about the future is always only more or less uncertain then Risk must be both related to and distinguished from Uncertainty merely by degrees of support for any particular estimate of the probability of outcomes. Risk and Uncertainty are not discrete knowledge situations but rather represent the idealised ends of a continuum of situations where an actor’s knowledge of the probability of events moves from less certain to more certain.

Risk was subsequently conceptualised as a product of inquiry. The stages of the idealised Risk Management process can be used as a heuristic device to illustrate this. That is, it is normatively expected that an actor's understanding of risk should evolve from an initial state of uncertainty, generality, ambiguity, and vagueness (i.e. risk is not well understood) to one of relative certainty, concreteness, clarity, and specificity (i.e. risk is well understood) as he or she progresses through the stages of the risk management process, and particularly through the risk assessment phase. In the beginning actors may be uncertain about what the relevant risks are, and may only be able to talk about them in vague and ambiguous terms. But by the time they have properly defined, analysed, and evaluated those risks, they should be certain about the relevant risks, be able to talk about them in clear and specific terms, and, ideally, be able to quantify the real probability distribution of outcomes.

However, the ways in which the term "risk" is practically employed in organisational contexts "muddies the clean lines of its conceptual specification", to borrow a turn of phrase from Castree (2006, p 1). At the most generic level, the term "risk" is used to refer to the inherent uncertainty of the world. That is, in any complex system the future cannot be predicted with certainty because there are always non-trivial possibilities for the occurrence of events and situations with unexpected outcomes. In this very general sense, the term "risk" is synonymous with uncertainty as an inherent part of being in the world; it expresses the notion that actors are exposed to risk all the time whether they are aware of it or not (Rescher 1983). Risk becomes specific, which is to say specifiable, when actors imagine and make explicit particular states of the world as more or less likely to come about in the future. The term "risk" now refers to a specific event or situation, with specific outcomes and a certain likelihood of realisation in the future, although the precise nature of the event, its outcomes, and the probability of their occurrence may remain more or less uncertain depending on the actor's state of knowledge (Rosa 1998; Thompson 1986). Thus, as a practical category, risk not only encompasses the full gamut of possibilities for loss or gain in the future, from the trivially mundane to the improbably significant, but also is employed whether the event in question is only a vague uncertainty or a precisely known probability. Certainly, when actors talked about risk at Watercare they were typically referring to a general category of perceptions about future states of the world, where that category was defined only by the conjunction of negativity and uncertainty, and only rarely to an objectively quantified estimate of probability.

Decision theorists might argue that such ambiguity in the practical usage of the term "risk" is a problem of loose terminology on behalf of the actors in the local context. That is, situations where uncertainty is quantifiable should be properly classified as cases of decision



making under risk, while all other situations involving non-quantifiable uncertainty should be classified as decision making under uncertainty (Ang and De Leon 2005; Aven 2004; Hansson 1994; Kleindorfer et al. 1993; Moore and Thomas 1975; see also Chapter 6); and the local actors should be educated to know the difference. But while theoreticians may bemoan careless terminology, the definitional ambiguity found in organisational discourse is also the product of practical constraints on the calculation of risk. In order to talk about something called risk in ways conforming to the narrow quantitative specification of that concept in theoretical discourse requires the production of highly objective and precise knowledges about the real natural variation of phenomena in the world. Such states of knowledge are producible only with substantial data sets and complex infrastructures of calculation that are difficult and costly to construct in practice (Callon 1999; Covello and Mumpower 1985; Latour 1987, 1988; Law 1992; Law and Singleton 2000; Rechar 1999). Consequently, such calculative circumstances are the exception rather than the rule, and in most organisational contexts statements of risk represent more or less subjective perceptions, in the Bayesian sense, of the world around us. In these contexts, risk is less a specific state of knowledge about the future and more a label for a category of perceptions of what might or could happen, both specifically and in general. Risks are imaginable possibilities (i.e. events, situations, scenarios) for the future, and actors may know more or less about the likelihood of those events occurring, the nature of the outcomes associated with them, and whether and to what degree those outcomes will be significant.

Individuals and organisations may identify (or otherwise imagine) possible future scenarios as risky, analyse and evaluate those scenarios as warranting organisational attention, and develop and implement mitigation strategies; “humans can and will make causal connections between actions (or events). Consequences are perceived... [and] can be altered either by modifying the initiating activity or event or by mitigating the impacts” (Hilgartner 1992; Hutter and Power 2005a; Renn 1998, p 51; Ward 2005). Organisations can thus be seen as engaged in a constant cycle of reorganising how they pay attention to uncertainty through ‘encounters’ with risk (i.e. errors, accidents, and anomalies) which throw into question the existing ‘organisation of attention’ and generate processes of sense-making about what happened, how, and why, and what to do about it (Hutter and Power 2005). These activities, which can be understood as the organisational equivalents of stimulus, information processing, and adaptation, are not isolated activities that the organisation performs in addition to everything else, but rather must be seen as a constitutive feature of organising activity (Hutter and Power 2005):

Management orthodoxy suggests that organizations represent co-operative endeavours which seek to process and manage different sources of uncertainty in the pursuit of a goal, e.g. profit. Accounting and information systems, strategic planning processes, human resource and marketing functions, regulatory compliance and procurement processes are all components of this management of uncertainty in its broadest sense.... the managing of risk in general is a constitutive feature of organization and is not some accidental feature of it. (source: Hutter and Power 2005, pp 1 – 2)

In this implicit sense the management of “risk” (risk management, *uncapitalised*) is what organisations and their officers and employees do *in general* to anticipate and manage the gamut of possibilities for loss or gain in an uncertain future. The CRM expressed this point in Dialogue 13 when he said that “Good practice, good management is already risk management” (p 327). In other words, managing “risk” is already core business, and not something extra that organisations have to start doing; and to assume otherwise is to make what Ryle (1949) called a category mistake:

A foreigner vesting Oxford or Cambridge for the first time is shown a number of colleges, libraries, playing fields, museums, scientific departments and administrative offices. He then asks ‘But where is the University? I have seen where the members of the Colleges live, where the Registrar works, where the scientists experiment and the rest. But I have not yet seen the University in which reside and work the members of your University.’ It has then to be explained to him that the University is not another collateral institution, some ulterior counterpart to the colleges, laboratories and offices which he has seen. The University is just the way in which all that he has already seen is *organized*. When they are seen and when their *co-ordination* is understood, the University has been seen. (*excerpt from Ryle 1949, p 16, emphasis added*)

The parallel here is that asking “where is the risk management?” assumes risk management to be a collateral process, when it is really just the way in which a firm’s processes for paying attention and responding to uncertainty are organised and co-ordinated. The distinction has implications with respect to how we should understand the risk management division of labour, i.e. (1) the taking of risk, undertaken by firm management, headed by the Chief Executive; (2) the observation and support of risk management, undertaken by the office of the Chief Risk Officer; and (3) the monitoring and audit of risk management, undertaken by the office of the Internal Auditor (Acharyya 2008). If, practically speaking, “risk” is synonymous with uncertainty, then the signifier “risk” becomes superfluous, and indeed distracting. The division of labour can instead be seen simply as: (1) management, which may be understood as the task of paying attention and responding to uncertainty, undertaken by a firm’s managers and employees, led by the CEO; (2) the observation and support of management, undertaken by the office of the CRO; and (3)

the monitoring and audit of management, undertaken by the office of the IA.

Somewhat confusingly, however, Risk Management is something that firms do in addition to everything else. Firms have Risk Managers or whole Risk Management departments, they have Risk Frameworks and processes and procedures for Risk Management, and stakeholders, managers, and employees all talk about Risk as something specific to be managed in addition to everything else that the firm does. In this sense, Risk functions as an explicit “organizing category” for management, a concept in whose name organising and re-organising activity is performed (Hutter and Power 2005a, p 9). But what then is the difference between those explicit instances of Risk Management performed in specific times and places within the organisation, and the various other ways in which organisations and actors pay attention to uncertainty in general? That is, what distinguishes “Risk Management” from “risk management”?

That Risk Management might be defined in relation to Risk can be dismissed. The point has already been made that Risk is a very specific and well developed form of knowledge about the future, which in most cases is both difficult and costly to produce, and requires specialised knowledge production capabilities. It has also already been demonstrated that Risk Management encompasses a broader category of knowledge production processes and activities, only some of which result in the production of properly specified and quantified statements of Risk consistent with the theoretical definition. Therefore, somewhat paradoxically, Risk is not a defining criterion of Risk Management.

The distinction is also not one of the nature of the activities in which actors might be engaged. Risk Management and risk management both encompass processes and activities for anticipating and responding to uncertainty. Although the generic designs for Risk Management depicted at the world-level might lead us to expect Risk Management to take on certain distinct forms (i.e. formality, explicitness, consistency, rigorousness, etc.), this is not necessarily the case. The activities in either category may take on a variety of forms, ranging from informal, ad hoc practices through to formal processes or procedures. The nature of the activities involved is therefore irrelevant to any distinction between risk management and Risk Management. Indeed, the same activity might be categorised as either risk management or Risk Management depending on the context in which it is performed.

If there is a certain identifiable characteristic which could be said to be unique to Risk Management, then it is merely that Risk Management is explicitly called “Risk Management” by the actors involved. That is, the actors are aware that they are performing some sort of Risk Management activity, they talk explicitly about risks and risk controls, and the activity is carried out under the umbrella of some formal Risk Management programme or framework.

In contrast, actors engaged in risk management do not think of themselves as engaged in performing Risk Management activities, and do not make explicit use of Risk Management terminology (or if they do its use is incidental). But if there is no substantive difference in the way these two categories of activities are performed, other than the explicit use of Risk Management terminology, then what purpose does Risk Management serve, and what does it achieve that is not achieved through risk management?

The answer to the first part of the question is normative, grounded in the emergence of Enterprise Risk Management in the 1990s as a codified solution for good governance, internal control, and decision making in organisations (see Chapters 1 and 3). The ultimate vision of ERM is for the comprehensive and systematic incorporation of uncertainty (i.e. consideration of threats and opportunities) in organisational decision making, through formal, rigorous, scientific processes, embedded throughout the enterprise, both structurally and culturally, and supported by specific tools and calculative infrastructure (MacGillivray et al. 2007a, 2007b; Pollard et al. 2004; Ward 2003; 2005; see the overview of the features of mature ERM in Appendix V). In this regard, the appointment of Risk Officers and Risk Committees, the development of Risk frameworks, registers, policies, and procedures, the explicit assessment and representation of Risk, and the formal reporting of Risk to the Board, among other activities, all demonstrate organisational commitment to the ideals of good governance and decision making (Power 2004, 2005b, 2007).

But to the extent that organisations already account for and manage uncertainty in a variety of ways, the various international frameworks and capability maturity models for ERM can be reinterpreted as more than just designs for a collateral Risk Management capability. More broadly they can be seen to constitute a standardised, though generic model for evaluating how the management of uncertainty is organised, co-ordinated and undertaken (Power 2004, 2007); in other words for evaluating the quality of risk management. The following excerpt from Ward (2003) illustrates the interpretation:

Every member of an organisation needs to make decisions, plan, and manage uncertainty to a greater or lesser extent, so RM should find natural application in all organisational activity. Indeed, much existing good management practice could be regarded as RM. For example, effective planning, coordination, setting objectives, and control procedures are all responses to pervasive sources of risk such as human error, omissions, communications and so on. Nevertheless, the extent and quality of RM carried out in an organisation can be very variable, ranging from sophisticated, formal processes in some areas to a reluctance to contemplate uncertainty in any form. A common intermediate approach involves informal processes, often involving little more than an intuitive perception of risk, followed by ad hoc approaches to the management of risks. In the absence of formal processes, RM is often implicit within existing business processes, so that it is less effective than it could be. The challenge is how to develop

RM practice in ways which increase the extent and effectiveness of RM in the organisation.  
(excerpt from Ward 2003, p 7)

Watercare's CRM expressed this interpretation in Dialogue 14, when he suggested that Risk Management was not "some kind of new process that the organisation is going to start doing" but was rather a "polarising lens" for looking at what the business already did (p 335). From this perspective, ERM frameworks and capability maturity models do not prescribe a procedure to be developed and followed, but rather present a standardised model against which to compare and evaluate the quality of the processes through which a firm manages risk (uncertainty). Although there is no reason to assume that those processes must be ad hoc, informal, unscientific, subjective, etc., simply because they have not been evaluated and developed under the auspices of a Risk Management programme, to the extent that those extant capabilities are revealed to be other than "advanced" or "mature", the ERM imperative is that they could, and should, be made more effective by developing them to a more advanced level. This is particularly evident in claims that ERM capability maturity models constitute normative guides for the development of risk management systems and processes (e.g. Pollard et al. 2004, p 460).

The critical assumptions of such a view are (i) that a more advanced or mature capability for managing uncertainty is more effective with respect to protecting and creating shareholder value than a novice capability; and (ii) that if moving from less mature to more mature requires a movement from the ad hoc, informal, implicit treatment of risk (uncertainty) to consistent, formal, and explicit modes of treatment, then the effective management of uncertainty necessarily exhibits the characteristics of mature ERM, as depicted in international capability maturity models. The all-encompassing nature of the ERM concept implies a further assumption, that (i) and (ii) apply with respect to any decision or management process in the organisation, from the routine to the strategic, suggesting a virtually unlimited purview for ERM and thus also for the Chief Risk Officer.

Enterprise Risk Management frameworks and capability maturity models are, however, normative rather than descriptive. They purport to define the characteristics of effective risk management, but those definitions are not necessarily based on sound empirical analysis of actual risk management processes in organisations. Rather, the process definitions and attributes of mature ERM appear to be derived primarily from a top-down translation of the principles of ERM, with the various models achieving this in more or less detail (see Appendix V). Whether, to what extent, and under which conditions the above assumptions actually hold true in organisations is a matter for research. But academic research on the firm-level effects of ERM implementation has to date been limited, demonstrating a range of

idiosyncratic approaches and inconsistent results, and none has specifically sought to correlate enterprise performance with risk management process characteristics (Iyer et al. 2010; see the review in Chapter 1). Thus, the answer to the second part of the above question – what does Risk Management achieve over risk management? – remains inconclusive.

### **The CRO's dilemma: how to convince actors to adopt more advanced capabilities**

Whether the above assumptions hold true or not, the interpretation of ERM as a “quality standard” for risk management is no less problematic for Chief Risk Officers to implement than the CRM's earlier interpretation of ERM as something to be “designed and built” in the organisation. As the champion of ERM, and therefore, according to the above assumptions, the promoter of more advanced capabilities, the CRO is faced with the challenge of convincing other actors to engage with and use different procedures, methods, or tools for paying attention to uncertainty. To do so foreseeably requires the CRO to demonstrate or argue for the following:

- i) That those procedures, methods, and tools will, or will likely, reveal new distinctions not currently perceived by local actors; and
- ii) That those distinctions would provide a better and more defensible basis for decision making and action than the status quo; and/or
- iii) That acting on the basis of those distinctions would be necessary or beneficial (i.e. performance of current capabilities would be more reliable, losses would be avoided, opportunities would be realised).

The methodological difficulties of making those predictions notwithstanding, they can also only be legitimately made in context of the specific knowledge production processes in question. That is, the legitimacy of any such claims depends on the specific needs of decision makers. When the task of justifying new capabilities is contextualised in this way, a dilemma becomes apparent. For Watercare's CRM, it was this: he lacked the legitimising knowledge, experience, and authority to reasonably convince organisational actors that they needed more detailed “risk data” to make better decisions. For CRO's more generally, the problem lies with the generalist nature of the role. The long tradition of risk management as both a practice and field of academic inquiry in various disciplines has produced a correspondingly large and well developed body of knowledge on various types of risk and their management. Organisations face myriad different types of risks and may draw on various disciplinary

perspectives and methodologies, both formal and informal, for their analysis and management. The particular approaches adopted in any given decision making context will depend on the nature of that context, the significance of the risks, and the purpose of the risk assessment. Modelling the reliability of system components in a water treatment plant, for instance, involves very different knowledges, analytical methodologies, and calculative infrastructures than brainstorming potential strategic issues affecting the enterprise as a whole. The background and experience of the individual CROs notwithstanding, they are unlikely to be experts in everything, and therefore will not have the legitimate knowledge and skills to make the above justifications in every context for which Risk Management oversight of risk management is required. Indeed, it is more likely that CROs will know less than the actors they are trying to convince. In the absence of the specific expertise and hence legitimacy needed to convince organisational actors that they should adopt new or more advanced capabilities, generic appeals to normative capability maturity models are unlikely to be sufficient. This is the general form of the earlier dilemma encountered by Watercare's CRM.

### **The potential to add value as a “strategic advisor”**

The CRM's second realisation opened up a much broader range of possibilities for the implementation of ERM and the performance of his role. In particular, for the CRM it emphasised that he now saw his role as, first and foremost, about facilitating insight, and only in some cases about generating data. This paradigm shift was complete by the time Dialogue 14 took place, and the CRM expressed it in those explicit terms in that discussion. The CRM summarised the transition as follows:

Yeah, it's definitely a 180 degree turn, from saying that I've got to generate good quality data to saying that I could do this job by not generating data, that all the job requires to be done efficiently is to have someone who's insightful. You need someone who can look at systems, patterns, interactions, and who can cut to the chase. But the irony is that there's still a part of me that says “but data quality's everything, data quality's everything”, and that must be the engineer. I suspect if you read this transcript against the first transcripts from way back in May or June last year, you will see back then that I thought things were so much more decisive, so much more analytical. As you say, I was talking about data quality for objectivity, and now I'm at the opposite end of the spectrum saying nothing's simple, it's all complicated, there's no right answer, so I guess you could say it's a complete paradigm shift. But the irony is I feel like I'm better focussed on the core purpose of risk management, which is really concerned with making the organisation less vulnerable, a process of continuous improvement. I think now I could do it a lot better by just going out,

talking to people, looking at the business, cutting it and slicing it a hundred different ways, until we see things that drop out as important. (Dialogue 14, pp 333-334)

In the context of Dialogues 11 – 14, the CRM's proposal that Risk Management was a "polarising lens" for interpreting the business in terms of risks and controls constituted a methodological paradigm shift from the radical pursuit of an idealised state of objectivity to a more pragmatic approach of facilitating organisational learning. Under the CRM's proposal formal risk assessment was no longer (or not just) a procedure for producing Risk (or "risk data"), as a particular, quantified form of knowledge about the future variability of system performance or decision outcomes, but was reconceptualised as a process through which organisational actor's might make explicit, reflect on, and develop their knowledge of business systems and processes, of decision options and consequences, and of roles, responsibilities, and professional practices:

There's an analogy that keeps rolling round in the back of my head, which I picked up at some sort of leadership training and management course that I went to in the past. Inevitably the focus of those courses is self-awareness and they had this phrase that I thought was incredibly corny. It was very American. They said that as a manager you have responsibility. And it was this phrase, you have responsibility, and they hammered it throughout this course. What they were saying was analyse that phrase, you have response-ability, which is, and their whole idea was, that as a manager you have the ability to choose how you respond, you need to be aware of that, you need to have a wider awareness, you need to have self-awareness. And so I think that aligns very well with... this idea of reflective thinking. When something comes up do you just react or do you think about how you're going to respond, what action you should take? So there are good parallels in terms of personal development and organisational development. That's all that I seem to be doing is trying to generate self-awareness for the organisation. (Dialogue 12, p 324)

The goal was no longer Risk, specifically and ideally, but knowledge, generally and pragmatically. Rather than generating data the CRM now believed that he could add more value by providing spaces and opportunities for organisational actors to reflect on their activities, by facilitating the sharing of information across different communities of practice, and by revealing and challenging existing assumptions. In this regard, I would characterise the CRM's reconceptualised role as being consistent with the "strategic advisor" archetype defined by Mikes (2010). CROs of the "strategic advisor" variety exhibited what Mikes referred to as *quantitative skepticism*, characterised by a belief that risk modelling is "not sufficiently accurate to produce an objective picture of the underlying risk profiles" and that such quantitative calculations should be used as trend indicators only (Mikes 2010, p 78). The



emphasis in the “strategic advisor” role is on playing “devil’s advocate” and facilitating the cross functional sharing of risk information to prevent “risk incubation” and to enhance “risk anticipation” and learning about risk profiles, particularly for non-quantifiable uncertainties (Mikes 2010). The CRM described his reconceptualised role in just these terms.

## **How should CROs support organisational decision making?**

Chapter 4 and this chapter have demonstrated that the decision support function of ERM can not be effectively fulfilled by focussing solely on the objectivity of information or calculative methods. So how then should CROs support organisational decision making? The CRM’s reconceptualisation of his role, above, seems productive. His success with the strategic risk assessment at Watercare suggests that CROs might productively employ a facilitative approach to such exercises. This was supported by the analysis in Chapter 4 which positioned such roles as important for assisting knowledge production and learning in organisations.

However, as was explained in Chapter 2, I left Watercare shortly after Dialogue 14 was recorded in May 2008 to concentrate on analysing and documenting the data I had already collected. While I did keep in occasional contact with the CRM after that, I did not formally collect any further information about his experiences. As a result this thesis does not report on whether the CRM’s “paradigm shift” did in fact prove productive with respect to developing methodologies to support decision making within the company. As it turned out, the CRM focussed primarily on developing the assurance function of ERM after my departure, and that focus continued when he moved to take up the position of Risk and Assurance Manager with another company in 2010.

There is a variety of ways in which the implications of the CRM’s reinterpretation of Risk Assessment could be explored vis-à-vis decision support. As I was following the CRM, my intuition was to explore the apparent similarities between the normative designs for Risk Management in the international literature and the traditional depiction of a rational decision in literature on decision making. Chapter 6 therefore examines, from a theoretical perspective, how people make decisions and the various factors which influence decision making in practice. That account is then used to frame up potential strategies by which Chief Risk Officers can support organisational decision making under an ERM framework.



# Chapter 6

## Framing up decision support strategies for Chief Risk Officers

Risk management is a means to an end. The main game is to make good decisions. The bottom line for leaders is that they must make decisions that lead to success for their organisation. How and why decisions are made is a field of study in its own right, but it is clear that consideration of threats and opportunities should be part of all good decision making. Nonetheless, it is the quality of decision making not the quality of risk management that lies at the core of success.

*Dr. Richard Barber (July 2010) Getting the assumptions right.*

*RiskPost – The Journal of the New Zealand Society for Risk Management, 10(2).*

Chapter 6 addresses from a theoretical perspective the question of how Chief Risk Officers might fulfil the decision support function of Enterprise Risk Management. The chapter takes as its starting point the following assumption, which is illustrated by the above quote: if Risk Management is fundamentally about making good decisions, and if the CRO is the agent tasked with organising and co-ordinating Risk Management enterprise-wide, then the strategies, methodologies, and tools that CROs employ should be grounded in a theoretical framework which accounts for how agents make decisions. In other words, if the objective is to help agents make better decisions, such that the CRO is, in this regard, a kind of Decision Engineer (March 1978), then he or she requires an understanding of what constitutes a good decision, how people make decisions in practice, and the various factors which influence their ability to do so well. The objective of Chapter 6 is to formulate such an understanding, in the form of an orienting framework, as a basis for conceptualising strategies by which CROs

may productively intervene to support decision making in organisations.

To avoid confusion, I treat “decision making” and “problem solving” as identical processes. A “decision” is commonly understood as the selection of a course of action from at least two alternatives, whereas a “problem” is usually understood to be some state of affairs which we would like to be otherwise, and which is more or less difficult to resolve (Buenaño 1999; Schermerhorn 1996). But problems always lead to decisions about what to do (even if only to preserve the status quo), and decisions always imply an imperative to change one’s circumstances. That is, problems can not be resolved without identifying and evaluating potential solutions, and decisions can not be made without first distinguishing what is problematic about the present. In this regard, the everyday notions of “problem” and “decision” merely refer to two idealised moments located at different ends of a common process. Whether we call that process “problem solving” or “decision making”, we are referring to the process by which humans perceive the world, formulate distinctions, and calculate courses of action.

The process and outcomes of decision making by individual actors, and by organisations, may be judged against various criteria, both normative and practical. The decision making process may, for instance, be judged against standards of objectivity, transparency, inclusiveness, and fairness, while the chosen option may be judged relative to alternatives in terms of appropriateness, effectiveness, and efficiency. There is, in this regard, a long tradition in various sciences of assuming that humans do or should behave in accordance with the precepts of rational choice theory (Cabantous, Gond, and Johnson-Cramer 2008; Hodgkinson and Starbuck 2008; Langley et al. 1995; Jaeger et al. 2001; Laroche 1995; Kleindorfer et al. 1993; March 1988b). This chapter upholds this tradition in the normative perspective in so far as it proceeds from the primary assumption that a “good” decision is, first and foremost, one that is rational (because if we are not rational in our endeavours then nothing else really matters; Rescher 1995).

The first section of the chapter briefly reviews the current standing of rational choice as a theory of human decision making in both descriptive and normative senses. It is concluded that more than 60 years of accumulated research has resoundingly challenged the canons of rational choice as a descriptive theory of how humans do actually make decisions, but that rationality cannot be abandoned as the normative standard for human decision making. This sets up the objective (and challenge) for the chapter to conceptualise and translate this ideal standard of “rational choice” into practical strategies for supporting decision making in different contexts within organisations.

The second section of the chapter compares contexts of skillful and deliberative action in

order to generate understanding about the agent-specific and context-specific factors which enable and constrain rational calculation. It is argued that those factors define a subjective relationship between the agent, who must calculate, and the domain in which those calculations must be performed, giving rise to different forms of decision making behaviour. Drawing on relevant literature, the section then collates that understanding into a generic typological framework of decision situations differentiated by: (i) the degree to which the agent perceives the problem domain to be uncertain and novel; (ii) the degree of difficulty experienced by the agent with respect to framing and calculating in that domain; (iii) the nature of the primary uncertainties with which the agent is concerned; (iv) the epistemological (knowledge) and axiological (values) dimensions of the framing task; (v) the nature and functions of the problem solver roles in each type of problem domain; and (vi) the translation of interests and the transmission of information which define the relationships between those roles and functions. The framework thus makes explicit key general features of the range of contexts in which a Decision Engineer might be required to intervene in order to help agents make better decisions.

It is clear from this framework that real decision making behaviour is only approximate, in form, to the traditional ideal of the rational choice in a limited set of circumstances; i.e. as an explicit procedure by which an agent identifies some preferred course of action which optimally satisfies all decision criteria and rules. This implies that it is necessary to disconnect the normative criteria of rationality from the assumed form of a decision. The third section of the chapter therefore reconceptualises the idea of a decision in performative terms. I borrow Callon's (1998a, 1998b, 1999) notion of "framing" as a process of disentangling, from everything else, what is relevant and significant for the task at hand, to explore the question of what it means to settle a decision frame under conditions of the rational imperative (i.e. that one must always do what is best). That is, if frames must be constructed as a series of distinctions about what is relevant and significant, then what are the conditions for settling decision frames rationally? The discussion does not try to define abstract parameters or conditions of rationality against which real decision making might be judged, since such a task would, I think, be a case of *reductio ad absurdum*. Rather, the analysis calls attention to those sources of uncertainty which are common to rational calculations everywhere, since helping actors to resolve judgements with respect to those uncertainties would be the goal of decision support interventions.

The last section of the chapter draws on the earlier work to frame up a general strategy for decision support by Chief Risk Officers. It is argued that, due to certain unique constraints on the CRO role, CROs should focus their attention, vis-à-vis the decision support function of

ERM, on one particular methodological approach (two other approaches are also identified, but are dismissed as unsuitable as *general* decision support strategies for CROs). The methodological approach is not defined in detail. Rather, the discussion focuses on identifying those features which are likely to be common to the approach across different contexts and those factors which are likely to differentiate the approach across contexts. In this regard, the chapter presents a strategic outline of the methodological approach which could serve as a stepping off point for future research and development in this area.

## The status of rational choice theory

The canons of rational choice theory portray decision making as intentional, consequential, and optimising: “they assume that decisions are based on preferences (e.g., wants, needs, values, goals, interests, subjective utilities) and expectations about outcomes associated with different alternative actions. And they assume that the best possible alternative (in terms of its consequences for a decision maker’s preferences) is chosen” (March 1988b, pp 1-2). The normative model of rational choice has traditionally been depicted as a series of consecutive stages in which the agent identifies and ranks, in order of preference, possible future states of the world, the various courses of action by which they might be achieved, and the expected outcomes of those actions (Callon 1999, p 190; Hansson 1994). That is, in order to calculate the rational choice in a particular situation, the agent must be able to do the following (modified from Callon 1998b, p 4):

- 1) Establish a list of the possible future states of the world (each state being defined by a certain arrangement of people and things);
- 2) Rank those states of the world in order of preference;
- 3) Identify the actions which allow for the production of those states of the world;
- 4) Identify the consequences of those actions; and
- 5) Evaluate the alternative courses of action against (2) to identify the best course of action.

When this is achieved, the rational choice is merely that course of action which is ranked at the top of the list. But since a rational choice is, by definition, only such because all the other options were considered to be less appropriate, perfect rationality would seemingly require the agent to *know* that there are *no other* possible options which might be better. In

other words, perfect rationality requires the decision situation to be fully specified (except for what might be called really indeterminate or irrelevant factors) and the correct methodological procedures to be correctly applied (March 1978). The decision maker must know *all* his or her objectives and priorities, identify *all* options, assess *all* consequences, have *all* the relevant information at hand, analyse *all* without error to identify the rational choice, and then must take it (Callon 1998b).

However, since the 1950s research in a number of fields has compellingly demonstrated that people often violate the tenets of rational choice theory (Cabantous et al. 2008; Hodgkinson and Starbuck 2008; Langley et al. 1995; Laroche 1995; Shafir and LeBoeuf 2002; Weber and Johnson 2009). A full review of the accumulated body of research on judgement and decision making is beyond the scope of this chapter, but it is clear that two fundamental assumptions of the traditional model have been called into question.

The first is the assumption of completeness. The idealised specification of the rational choice is predicated on a “view-from-everywhere, which “presupposes the existence in organized form of all the relevant information on the different states of the world and on the consequences of all conceivable courses of action and the access of all this information to the agent” (Callon 1998b, p 4). But such a synoptic procedure “cannot be practiced except for relatively simple problems and even then only in a somewhat modified form” (Lindblom 1959, p 80). In most real-world contexts, decision-making confronts the difficulty and expense of the framing process, usually in the form of limits imposed by time and resource constraints and the cognitive capabilities of the actor’s involved, such that it is rarely, if ever, possible for actors to perform a rational calculation in a form consistent with the idealised normative specification (Lindblom 1959, 1979; Loasby 1999). Rather, practical approaches to decision making tend to be incremental and marked by the use of a “mutually supporting set of simplifying and focussing strategems” (Lindblom 1979, p 517). Weber and Johnson note as follows (2009, pp 76-77):

Although we are restricted by finite attentional and processing capacity, we also are blessed by an abundance of ways in which we can focus and utilize this finite capacity... We apply a wide repertoire of processing modes and strategies to our choices and inferences in a fashion that is cognizant of our goals, capacities, and internal and external constraints... some decision strategies are more automatic, associative, and affect laden, whereas others involve either implicit or explicit attempts to consider the pros and cons of different choice alternatives.

Thus, people rarely, if ever, calculate in an explicitly rational sense, but instead “muddle through” using a range of strategies to simplify the calculative task and cope with their information processing limitations (Dane and Pratt 2007; Hodgkinson and Starbuck 2008;

Lindblom 1959, 1979; Shafir and LeBoeuf 2002; Weber and Johnson 2009).

The second core assumption of rational choice theory which has been compellingly challenged is that of consistency. According to the normative model, the specification of what is rational or appropriate with respect to belief or action is contingent on the prior definition of certain objectives and preferences. That is, one cannot evaluate alternative possibilities unless one knows on what terms the comparison is to be made. Since people are entitled to a wide variety of opinions, beliefs, and preferences, it follows that what is rational is less to do with the substance or content of beliefs and intentions than with how agents come to hold them (Shafir and LeBoeuf 2002). Indeed, if rationality were solely a question of the substance or content of beliefs then this would imply that humans could be approximately ideally rational only in relatively simple situations where they could fully specify, value, and relate everything that needed to be taken into account in their calculations. In increasingly complicated situations the decisions performed by actors would, by definition, become increasingly less rational because they would be uncertain, in some significant sense, about the elements of the decision and of the problem domain. That is, the rationality of agents would, somewhat perversely, depend on factors beyond their control, such as the complexity of the situation, and the information and resources that were available to them. For this reason, rationality must be seen as a characteristic or quality of the decision maker's reasoning process, rather than of the outcome of that process (Shafir and LeBoeuf 2002). It is, in this sense, a person-relative concept. There is no external standpoint from which to evaluate rationality except that of consistency. An agent's beliefs and intentions should not obviously conflict with the agent's knowledge or preferences – they “should cohere, should adhere to basic rules of logic and probability theory, and should not be formed or changed based on immaterial factors related to, for example, mood, context, or mode of presentation” (Shafir and LeBoeuf 2002, p 493). However, “[m]any studies... have documented numerous ways in which judgments and decisions do not cohere, do not follow basic principles of logic and probability, and depend systematically on just such irrelevant factors. People use intuitive strategies and simple heuristics that are reasonably effective some of the time but that also produce biases and lead to systematic error” (Shafir and LeBoeuf 2002, p 493; see also Dane and Pratt 2007; Hodgkinson and Starbuck 2008; Hodgkinson et al. 2009; Weber and Johnson 2009).

The accumulated research which shows people systematically violating fundamental normative principles of reasoning constitutes a critique of rational choice theory which “is compelling and rightfully gaining influence in the social sciences in general” (Shafir and LeBoeuf 2002, p 491). That critique is generally targeted at what Shafir and LeBoeuf referred



to as “the rationality assumption” – the expectation that people are able to, and do actually fulfil or at least approximate the requirements of rationality in their endeavours – which “has come to constitute perhaps the most common and pivotal assumption underlying theoretical accounts of human behavior in various disciplines” (2002, p 492). In other words, the rationality critique challenges the veracity of rational choice theory as a descriptive theory of how people actually behave when they make decisions.

The fact that humans often fail to live up to the rationality assumption certainly raises an important question about the normative relevance of rational choice theory. That is, if the idealised conditions of rationality are realistically unattainable in practice, and if people frequently violate the tenets of rationality in the ways they “muddle through” the decision making process, then why should those conditions and tenets continue to be held as the normative standard for good decision making? Perhaps a more pragmatic standard is required, one that more accurately reflects how people make decisions in practice? However, as Rescher (1995) has pointed out, the imperative to reason cannot reasonably be abandoned. To eschew the rational imperative would imply that one no longer aspires to seek out and do “what is appropriate in matters of belief, evaluation, and action”, and although such a choice can be made, it can not be made reasonably (Rescher 1995). If the aspiration to reason is abandoned, then there is no other more reasonable alternative which could take its place (Rescher 1995).

That people frequently violate the conditions and tenets of rationality in practice is precisely the relevance of rational choice theory as a normative standard for how decisions ought to be taken (Jaeger et al. 2001; Kleindorfer et al. 1993; March 1978, 1988b). Indeed, if the conditions of rational choice were such that nearly everyone could fulfil them most of the time, then rationality would be trivialised as a standard for judgement and decision making, and there would be little to debate. From this perspective rationality is no longer some inherently human trait that agents bring “ready made” to the table when they want to make a decision, but is instead held as a normative standard – both abstract and ideal – which humans should aspire to, but realistically will only be able to approximate under the right conditions (Cabantous et al. 2008). Thus, contrary to our intuition, we should not take rationality for granted: even if we assume the rational intent of human actors, we should not assume that such intent translates easily into practice (Cabantous et al. 2008). This provides the motive for the decision support industry. Real decisions by human actors are either approximately rational (“good” decisions), or hopelessly irrational (“bad” decisions), or somewhere in between, and this means that there is, in principle, always room for improvement. Actors can, and should, be encouraged and helped to calculate “more

rationally”. The relevant question then is how to conceptualise and translate this ideal standard of “rational choice” into practical strategies for supporting decision making in different contexts within organisations.

## **Understanding real decision making behaviour**

The following sections compare contexts of skillful and deliberative action in order to generate understanding about the agent-specific and context-specific factors which enable and constrain rational calculation, and how those factors constitute a subjective relationship between the agent, who must calculate, and the domain in which those calculations must be performed, giving rise to different forms of decision making behaviour. Drawing on relevant literature, that understanding is then collated into a generic typological framework of decision situations, defining the general features of the range of contexts in which a Decision Engineer might be required to intervene in order to help agents make better decisions.

In this section, the term calculate should be understood as loosely referring to any process of relating, comparing, transforming, or manipulating information according to certain rules or procedures in order to draw a conclusion or judgement (Callon and Muniesa 2005). Such a definition not only encompasses operations in a mathematical or numerical sense, but also those less formal or rigorous forms of reckoning from means to ends (Callon and Muniesa 2005). The important part of the definition is that of comparing or manipulating things on the basis of some common operating principle, not the specific form the calculation takes.

### **Case 1: the formalised construction of a rational choice**

During the period of the research, the largest capital project then being planned by Watercare was the Hunua No. 4 watermain, a 1200–1900mm diameter pipeline to be constructed between central Auckland City and a major reservoir complex approximately 30km southeast of the city, with an initial estimated capital cost of ~\$NZ200 million. The main was required to provide for future growth in demand, and to mitigate security of supply risks carried on the other three Hunua mains. The planning process was exhaustive, involving the investigation and assessment of 52 route variations through western, central, and eastern parts of the Auckland isthmus. For each of those variations the potential route had to be defined to minimise construction costs, risks, and duration, and the impacts on the community and key stakeholders. Each option was modelled in the context of the expected

future (2057)<sup>14</sup> water supply network to identify the hydraulic impacts of the new pipeline and any additional works that would be required to optimise the operation of the network. Finally, each option was evaluated against 31 decision criteria (multi-criteria analysis). Through this process the option pool was gradually refined from the initial matrix of 52 options to 26, then 12, then to the final three options presented to Watercare's management, Chief Executive, and Board of Directors in the form of a Capital Expenditure Request (CapEx). That document made explicit the key components of the decision:

- The business need (in this case the need to address security of supply risks and provide head room to meet future demand growth in the city);
- The options for addressing that need (in this case several route options);
- The evaluation of those options (in terms of the financial costs, total risk reduction achieved, and the risk exposures of each option); and
- The recommended option.

Some characteristics of this “decision” making process may be noted. First, while the actual performance of the decision process was messy and not entirely linear, it was nevertheless a visible process which an observer could not only follow but also classify certain phases of activity as consistent with the stages of the idealised decision process (e.g. problem perception, data collection, option development, analysis, etc.). People could be seen performing certain tasks associated with the process and there were tangible records which could be accessed both before and after the process had been completed (i.e. emails, spreadsheets, models, reports, meeting minutes, presentations, maps, plans, etc.). And, the relevant facts, information, and decision elements (objectives, options, consequences, etc.) were fully made explicit in representational form, most obviously in the final report to Watercare's Board, such that they were visible to all decision makers. In this case, that report constituted a calculative space, being literally a place where the decision elements (or at least their representations) were identified, specified, and arranged, and which enabled the final calculation to be performed (i.e. the decision whether to proceed or not to proceed). Finally, there appeared to be an explicit moment and space of decision when the Capex request was presented to the Board of directors for discussion and approval (this being formally complete

---

<sup>14</sup> The planning horizon for the project was 50 years. The model of the 2057 water reticulation network included all existing infrastructure as well as planned additions and modifications that would be implemented before 2057. This provided the design operational context for the analysis of the new pipeline.

when the required signatures were affixed to the Capex document). This final moment of decision was, of course, ritualised. The actual conclusion about how to proceed may have been made some time previously, and indeed was probably reached by a number of different actors at different times, but a particular time and place can be identified where that conclusion was formalised by the affixing of signatures to a document. Prior to that act the “decision” was incomplete – it was simply an investigation with a conclusion. The distinction is that the conclusion could not be acted upon without formal signoff from decision makers, at least within the confines of the rules and procedures of the organisation, since it was the act of signing the Capex which formally conferred the authority to proceed and enabled the release of funds for expenditure.

Second, the decision making process was lengthy (taking longer than 12 months), and resource intensive. Watercare’s planners had to mobilise a metrological framework of instruments and devices, both physical and symbolic (representational), which enabled the identification, measurement, description, classification, and representation of the various decision elements and their relationships (Callon 1998a; Callon and Muniesa 2005; Latour 1987). In this case, the various investigative and calculative resources included (among others) route plans, geotechnical information, field visits, materials and construction cost models, construction schedules, hydraulic network models, community consultations, and risk analyses.

Third, the process also involved many people, representing diverse preferences, interests, and opinions; including engineers, technicians, consultants, managers, Board members, and community and stakeholder representatives. In this regard, the framing of the decision was not simply a matter of technical design, but was very much a political matter<sup>15</sup> involving concerns over disruptions to neighbourhoods and traffic flows, potential damage to infrastructure, services, and the environment, as well as, for instance, questions over how to balance the urgency of the project against the impact on Watercare’s capital expenditure budget, and whether and how to co-ordinate the laying of the pipeline with other major construction projects in the region, and even with Watercare’s own operational priorities.

The making of the Hunua No. 4 decision, at least as I have briefly summarised it here, is seemingly easy to characterise as the formalised construction of a singular rational choice. The example illustrates quite clearly the explicit and exhaustive (at least for all practical

---

<sup>15</sup> Judgements of relevance and significance, and therefore the inclusion or exclusion of any particular factor or entity from the decision frame, often depend as much upon political negotiation and agreement between actors as on demonstrating physical associations (Callon 1998a).

purposes) framing of the decision elements. It also highlights the mobilisation of an extensive calculative infrastructure of physical and symbolic devices which made that framing possible, and at least hints at the political nature of the process. While the magnitude of the investment in framing the decision is not immediately quantifiable, it is nevertheless evident in the materials, information, and people assembled in support of the Capex application to Watercare's Board. That assemblage represented the expenditure of considerable man-hours and technological, administrative, and political resources employed to establish the scope of the decision, to perform the necessary investigations and analyses to parse 52 route possibilities down to a final preferred option, and to bring that option in front of Watercare's Board as a recommendation agreed upon by many actors.

## Case 2: putting on a skillful performance

The following is an excerpt from A. J. Baime's article *War of Speed* describing the battle between Ford and Ferrari at the 1964 Le Mans (2009, pp 97 - 98). Driver Phil Hill in a Ford GT40 has just rejoined the race after a lengthy pit stop:

...By this time Hill was in 44th place. He'd lost 22 minutes. To catch up to the Ferraris from that distance would require the powers of a superhero. Hill knew this circuit better than any man... [and] began to rip off a series of perfect laps. Experience told him how to make up time at high speed without overtaxing the engine. There can be only one shortest distance around a racetrack, achieved when the driver chooses the perfect line through every turn. As Hill moved the car through a bend, he could ease the tires within an inch of the edge of the pavement.

In large part the race was won or lost on the rev counter, the rpm gauge staring the driver in the face from the center of the instrument panel. If Hill aimed to take a turn at 4,500 rpm, 4,400 rpm wasn't good enough. The difference between a four-minute lap and a 3:58 lap on this circuit equaled roughly 25 miles at the finish.

Fans watched Hill shriek down the pit straight. Thumbs clicked on stopwatches when he flew past the start-finish. He was cruising at 185 mph in fourth gear at 5,700 rpm. A slightly inclining right bend led him under the Dunlop bridge. He eased up on the gas, then accelerated again, shooting down a slope at 183 mph into the Esses. He downshifted to third, then second. Easy on the downshifts; no stress on the gear teeth or clutch plate. Hill left the Esses in second gear at 5,800 rpm – 82 mph. A hard brake down to 65 mph, a tight right turn onto the Mulsanne Straight, and he hammered the throttle. Third, fourth. The g-forces pinned him against his seat. A glance at the tach: 6,100 rpm. Two hundred mph summoned with his toe....

Hard on the accelerator. Second, third, past the signaling pits on the right, back up to 180 mph. Hill hurled the car through turns, rear wheels struggling for grip. The grandstands appeared in the distance. Hill gunned through that chasm, a huge valley lush with human bodies. Thousands of eyes followed the blue-and-white streak as it passed, a Ford car hurtling 185 mph on four patches of rubber.

No two laps were the same. Hill's brain filtered stimuli, automatically ranking them in

order of importance in nanoseconds. Photographers leaning in and waving at him. Pit signals: P2 (pit in two laps), P1, along with lap times. With each lap, fuel burned off, lightening the car, increasing its speed. His perception was near extrasensory. "True concentration is not aware of itself," Hill would explain. "The flagmen, unless they are holding a yellow flag or some such thing, are perceived and forgotten," Hill said. "A car you are overtaking is registered and erased as you safely pass."...

The above passage provides a reasonably detailed account of the performance and experience of driving a high performance race car. Indeed it almost qualifies as a thick description (Geertz 1973), the richness of the narrative being sufficient to virtually transport the reader into the cockpit with the driver. It is a description not just of any performance, but of a skillful one, as may be judged by the calibre of that performance (Ryle 1949). The driver does not simply keep his car on the track, but successfully negotiates "a series of perfect laps" at speeds up to 200 mph, choosing "the perfect line through every turn", easing "the tires within an inch of the edge of the pavement", balancing engine rpm against wear and tear, and shifting carefully to minimise stress on the gear box<sup>16</sup>.

It may reasonably be assumed that the driver is in the race to win it, or at least place as best he can, from which we can infer a consistent motive for his actions. The driver is also observed to consistently act in a way that is effective with regard to fulfilling that motive. At any given instant in the race the driver is faced with a range of possible courses of action: steering (*when? which direction? how much?*), acceleration (*when? how fast?*), braking (*when? how hard?*), shifting gear (*when? which?*). He also has available a range of information on which to base his actions, some of which are explicit in Baime's narrative: the rev counter, tachometer, and other instruments and gauges on the dash board, visual inputs about what is going on around the car, the noise of the engine, and, perhaps most importantly, the movement of the car and how it feels as it hurtles around the track. Then, in the next instant, the driver is performing certain of those actions: steering (*this way, just now, just this much*), acceleration (*just now, just this fast*), braking (*just now, just this hard*), shifting gear (*now, down*). Since which actions he performs and when are neither inevitable (i.e. he could always have done something else) nor the product of chance, the driver is clearly resolving possibilities for action. But the manner in which he is doing so is very different from the decision performed by Watercare's planners in the previous example.

In that case the planners and decision makers visibly engaged in a prior process of

---

<sup>16</sup> The driving team of Phil Hill & Bruce McClaren did not win the 1964 Le Mans, but by the time they were forced to retire with gear box problems they had moved up the field from 44th to 5th overall, with Hill setting a lap record in the process (Baime 2009).

deliberation, and that process involved explicit representation of the decision elements in the form of physical inscriptions. The various components of the decision were not only identified, taken into account, and related, but the components and their relationships were explicitly inscribed in material forms, such that they became visible and accessible to all decision makers (and to us as observers). Such explicit representation implies that the planners and decision makers were consciously aware of the decision elements they were taking account of (i.e. they had to be in order to communicate them in explicit forms). In the case of the race car driver there is no such explicit arrangement of the decision components within a physical calculative space. Certain bits of information (e.g. speed, fuel level, engine temperature) are made explicit within the physical space of the car cockpit (by various dials and gauges), but the driver's process of reasoning from that information to action is invisible to us as observers. We can only see the actions that he performs (where he looks, and the movements of his hands and feet). If the driver is framing and calculating then, to use a colloquialism, he apparently does so solely "in his head".

But at 200 mph the driver is clearly not first deliberating about what he is going to do, at least in a fully conscious sense, before he does it. The driver may be consciously focussed on the performance at hand, but he is only partially aware of all the things that he is doing, and of all the things that he is taking into account in order to put on that performance. Just as I do not have to consciously think how to type this sentence, the driver does not have to consciously think how to drive his car. He knows, for instance, when to brake, down shift, and turn in in order to enter a particular corner at the right time and speed, but he does not have to consciously acknowledge and process all the elements of that decision, or explicitly think how to perform each of those actions. This is not to say that the driver is unaware of making the judgement, any more than I am unaware of what I am typing. As he hurtles down the straight his attention is focussed on the approaching corner, the distance to it, the speed of his approach, and the location of the car on the track. But the cognitive process of perceiving these things, of ordering and relating them, and of arriving at the judgement to brake now, down shift now, and turn now is not a conscious deliberative process. There is no "ghost in the machine"; the driver's actions are not preceded by his intellectual acknowledgement of the various rules or criteria which govern those actions (Ryle 1949). The driver simply (although there is little simple about it) knows or feels the moment for action, just as he knows how to perform all the necessary actions (braking, clutching, shifting gears, steering) in the right order and timing; and this "knowing how" is only partially a conscious activity (Ryle 1949).

The driver's "know how" is not some innate skill with which he was born, but rather is a

capability acquired through long experience driving race cars. He can put on such a performance because he is highly trained, skilled, and experienced at the task of racing. It is in this regard that the driver's performance is no less difficult or costly to put on than the construction of the decision over which route to select for the watermain. The difference is that, as observers of the race, we do not see the prior investments which made the performance possible. The race car itself, for instance, represents a considerable investment in man-hours and technological and administrative resources, all of which is targeted to producing and maintaining a physical context which supports the driver's decision making on the track. Similarly, the driver himself represents considerable investments of time and resources (both his own and by others) to develop the skills and experience necessary to put on a world-class performance on the track. In this regard, the driver's performance in the race was only the last act in a process which, in all likelihood, was longer, more complex, and more costly than that involved with planning the new watermain.

### **Factors affecting framing and calculative processes**

Recent psychological models have distinguished between two systems of reasoning: a rapid, automatic and effortless, associative, intuitive process, that is beyond conscious control and which enables people to parse large amounts of information rapidly (System 1, intuition), and a slower, rule-governed, deliberate, and effortful process, which is conscious and analytically detailed (System 2, reasoning) (Chaiken and Trope 1999; Evans 2007, 2008; Gilovich, Griffith, and Kahneman 2002; Kahneman 2003; Weber and Johnson 2009). There is debate about the extent and ways in which the two systems interact, but System 2 is generally thought to occupy a supervisory role relative to System 1, "because System 2 knows the analytic rules that the intuitive System 1 is prone to violate and thus can intervene to correct erroneous intuitive judgments" (Weber and Johnson 2009, p 67; Kahneman 2003). System 1 is not restricted solely to processing inputs from the perceptual system, but also deals with concepts. The intuitive operations of System 1 generate "*impressions* of the attributes of objects of perception and thought" (Kahneman 2003, p 699, emphasis in original). Intuition refers to judgements based solely on these impressions. System 2, on the other hand, modifies those impressions to generate judgements which are reasoned, intentional, and explicit, even if they are not overtly expressed (Kahneman 2003, p 699).

The two cases discussed above would appear to be illustrative of these two different modes of information processing. That is, the decision constructed by Watercare's planners may be seen as a formalised case of explicit reasoning (System 2), while the skillful



performance put on by the race car driver clearly involved a mixture of both intuition (System 1) and implicit reasoning (System 2). But if Watercare’s planners and the race car driver were all doing the same thing (i.e. identifying, taking into account, and relating various bits of information in order to calculate a course of action), and if the planners and the race car driver were all experienced and knowledgeable individuals, then why were the two processes so characteristically different? Why was the driver – an expert in racing cars – able to decide things rapidly and repeatedly inside his head, so to speak, without explicit prior reflection, while Watercare’s planners – experts in asset planning – had to frame and reflect on their decision in such an explicit manner? The answer to this question lies in understanding the differences between the two practice contexts (Table 6.1).

*Table 6.1. Comparing two decision contexts and performances*

Planning a water main		Driving a race car	
Characteristics of the decision process	Contextual features	Characteristics of the decision process	Contextual features
Type: explicit reasoning (System 2)	Actors experienced with the <i>general type</i> of problem domain and calculations that had to be performed.	Type: mixture of intuition (System 1) and implicit reasoning (System 2)	Actor experienced with the <i>specific</i> problem domain and the specific calculations that had to be performed.
Explicitation: decision elements cognitively and representationally explicit. Decision aided by external infrastructure.	Initially high uncertainty over decision elements. Considerable work required to construct decision frame.	Explicitation: some decision elements cognitively explicit. Decision completely internal (i.e. in the driver’s head).	Low uncertainty over decision elements. Decision frame virtually complete.
Consciousness: actors consciously aware of decision elements and of the calculation.	Complexity of the calculation exceeded the cognitive capabilities of the actors involved (needed calculative aids)	Consciousness: driver only partially aware of all that he is doing to perform skilfully.	Complexity of the calculation did not exceed the cognitive capabilities of the actor involved.
Duration: slow		Duration: fast	

Watercare’s planners and the race car driver certainly both possessed the necessary “know-how” to carry out the particular calculations in which they were engaged. This much may be inferred from their respective performances, which were not those of incompetents or novices but of skilled and experienced professionals. Such “know-how” is not innate, but is rather learned through both classroom instruction and practice. Watercare’s planners, for instance, acquired theoretical understanding and expertise through university education (e.g. in Civil Engineering), and practical experience in designing water supply infrastructure through their individual careers. Similarly, the race car driver was not only well versed on the

finer points of race tactics and vehicle handling, but also well practised. Thus, to say that the planners and the race car driver knew how to perform their calculations is to say that they possessed the requisite prior knowledge to do so. Strictly speaking it is problematic to talk about knowledge as a possession, but the term will suffice here to convey the notion that Watercare's planners and the race car driver were able to draw on certain acquired theoretical understandings and practical experiences or skills relevant to their respective problem domains (Tsoukas and Vladimirova 2001). In this regard, both the planners and the race car driver were familiar, to a certain extent, with the problem domains in which they were acting, and with the calculations that they were performing.

However, while Watercare's planners and the race car driver were all professionals engaged in tasks for which they were well trained and experienced, the contexts in which they were acting were very different. At the beginning of their project, Watercare's planners were faced with a problem situation that was only minimally framed. The controversy over whether the new pipeline was actually required, and by when, may have been settled, but the rest of the decision frame had yet to be constructed. The precise objectives for the new pipeline were largely unknown, as were the evaluation criteria and their rankings, and the potential route, materials, and construction options. The planners also had little information about, for instance, ground conditions, materials and construction costs, stakeholder perspectives and preferences, the hydraulic impacts of a new pipeline on the existing and future water supply network, or the risks associated with the construction of such a pipeline. They were therefore faced with the substantial task of framing these and many other elements, and of constructing the necessary infrastructure to enable that framing (e.g. computer models, spreadsheets, reports etc.). In this regard, any familiarity that Watercare's planners had with the task of planning and designing a watermain was only very general. They may have planned other watermains, and were thus familiar with the overarching process, but this particular watermain was different in myriad ways, thus requiring the construction of the decision frame anew.

In contrast, when the race car driver stepped into the cockpit of the car he was immediately located in a highly defined decision making context. Not only was all of the information required to perform the necessary calculations already at hand (e.g. the tachometer, rev counter, feel of the car, etc.), but the driver's decision frame was already largely complete. In fact, because the driver was familiar with the task of racing in a very specific sense, i.e. specifically the task of racing that particular car on that particular track, he brought that frame virtually ready-made to the cockpit. Even before he climbed into the car, the driver was prepared to make the necessary calculations, almost as if he was bringing with

him a pre-formatted template into which he could simply plug the relevant information. Watercare's planners, on the other hand, had to first construct their template before they could proceed to the calculation.

Finally, the last significant difference between the two cases is that of the complexity, and hence difficulty, of the respective calculations performed by the planners and the race car driver. While the task of performing in a Le Mans race is far from easy or simple, even for an experienced driver, it is a fact of the case that given the requisite experience and skill the necessary calculations can be performed, quite literally, in one's head. As was clearly demonstrated by his performance, the driver was able to repeatedly and consistently calculate appropriate courses of action without the need for prior explicit deliberation, or the support of external framing infrastructure. In contrast, the final calculation performed by Watercare's planners was sufficiently complex that it could not be performed "in the head" of any particular actor: evaluating 52 route options against 31 decision criteria is simply beyond the cognitive capabilities of any individual human. A calculation of such complexity can only be performed with the aid of external infrastructure.

The differences in decision form between the two cases may therefore be attributed to differences in a number of factors which influence the ability of agents to frame a situation and calculate a course of action in any given context. At least ten such factors may be inferred from the analysis of the two cases above (these are listed in Box 6.1: Nos 1 – 8, 12 & 13). In addition, different agents, or groups of agents, may also exhibit differences in their decision making processes and outcomes according to "chronic differences in values and goals, presumably related to historical, geographic, or biological determinants, that focus attention on different features of the task environment and its opportunities and constraints" (Weber and Johnson 2009, p 72). For instance, research has demonstrated differences in decision making due to differences in gender, age, personality, numeracy (the ability to process basic mathematical and probabilistic concepts), and goal preference (e.g. preference for maximising versus satisficing versus regret minimisation) (Weber and Johnson 2009, p 72). The emotional state of actors can also be a significant variable affecting what actors focus on, how they interpret and value information, and as a motivator of action (Weber and Johnson 2009). Four further factors influencing the ability of actors to frame and calculate in a given situation may therefore be added to the ten already identified (these are listed in Box 6.1: Nos 9, 10, 11 & 14).

**Box 6.1.** *Factors influencing the ability of actors to frame and calculate in context*

1. The degree to which actors have access to information about the problem situation;
2. The investigative, metrological, and calculative resources available to the actors;
3. The actor's prior experience with performing the same or similar calculations in same or similar situations elsewhere, i.e. relevant skills or "know-how";
4. The actor's knowledge of theory, methods, and procedures relevant to performing the calculations;
5. The cognitive capabilities of the actors;
6. The complexity of the calculations;
7. The time available to perform the calculation;
8. The motivation of individual actors to pursue the calculation;
9. The actors values, whether personal, cultural, religious, moral, etc;
10. The personal biological and cultural biases and tendencies of the actors;
11. The emotional experience and mood state of the actors;

*And, where more than one actor is involved:*

12. The efficiency and effectiveness of communication between those actors; and
13. The ability of political actors to reach agreement on value judgements;
14. The cultural and emotional dynamics of the group.

## The relationship between agent and context

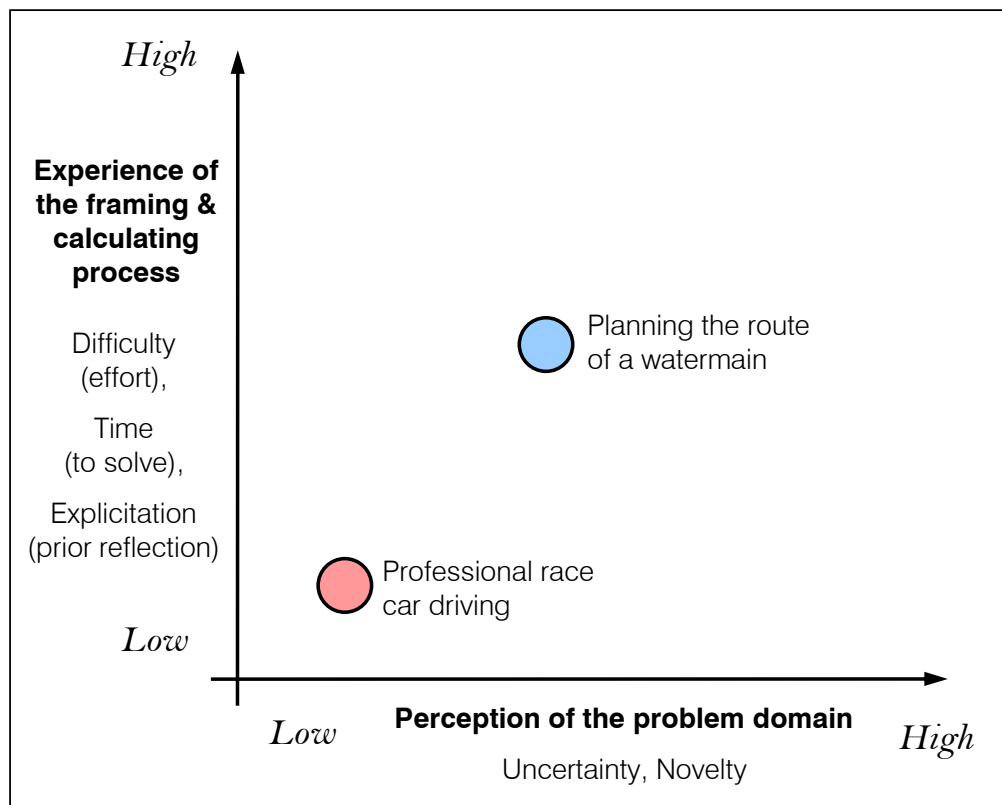
Some of the factors in Box 6.1 are skills, capabilities, or characteristics of the decision maker, while others are characteristics of the decision maker's context. Collectively they define a subjective relationship between the agent and the problem domain in which the agent is located. This means that whether a situation is problematic or not, and the nature of that problem, depends as much on who perceives it as such, and subsequently tries to resolve it, as it does on the objective characteristics of the situation (Buenaño 1999). In the terms of this discussion, the ability of an actor to frame the elements of a calculation, and subsequently to perform that calculation, is a function of both the actor's capabilities and the particular context in which that actor is located. It is not necessary to attempt to map out the full spectrum of combinations, but on the basis of the foregoing discussion it seems reasonable to infer that the factors in Box 6.1 will collectively influence both how an actor might perceive a given problem situation and that actor's experience of the subsequent framing process:

- **Actors' perception of the problem situation:** Some of the factors will influence how

an actor will perceive a given situation in terms of the degree of *uncertainty* (the extent to which the decision elements are unknown or unknowable by the actor), and *novelty* (the extent to which the actor is unfamiliar with the situation and the calculations he or she has to perform).

- **Actors' experience of the framing process:** Some of the factors will influence how that actor subsequently experiences the corresponding problem-solving/decision-making process, in terms of the *difficulty* of that process (i.e. the effort required), the *duration* (how long it takes the actor to complete it), and the degree of *explicitation* involved (i.e. the use of calculative infrastructure and the degree of prior conscious reflection involved).

Figure 6.1 conceptualises this relationship as a positive correlation between the degree of uncertainty and novelty perceived by the agent (x-axis), and the degree of difficulty experienced by that agent in framing and solving the problem, the duration of that process, and the degree of explicitation involved (y-axis). Both axes are therefore subjective to the agent. The two cases discussed earlier are differentiated on Figure 6.1 according to their relative characteristics in the two dimensions (cf. Table 6.1).



**Figure 6.1.** The relationship between agent and context with respect to framing. The x-axis represents the problem situation, as it is perceived by the agent, and the y-axis the nature of the problem solving process, as it is experienced by the agent. Both axes are therefore subjective to the agent.

The relationship depicted in Figure 6.1 highlights a significant feature of real-world decision making: that beliefs and intentions only become matters for conscious reflection when there is some uncertainty or difficulty with the underlying calculation (Rescher 2005; see the discussion on the pragmatist conception of knowledge in Chapter 2). In situations of near complete certainty, where the agent already knows what to do and how to do it, the agent's state of belief or intent is typically taken for granted. That is, the agent takes as given the objectivity of the knowledge at his disposal and moves easily, even automatically, from belief to action. Conversely, in situations of significant uncertainty and novelty, the agent cannot automatically assume the objectivity of the knowledge at hand, and faces a significant gap between belief and action. In these situations, the automatic flow of action is disrupted and the question of what we should believe or intend to do is forced to the forefront of our conscious reflection; the calculative process becomes increasingly conscious and explicit.

In the case of the race car driver much of his performance was automatic in the sense just described. He could move rapidly and repeatedly from perception to action because, for him, there was virtually no uncertainty or difficulty involved. His brain and body knew both what to do in response to those perceptions and how to do it. This is not to say that the driver performed unconsciously, but rather that long experience had provided him with the ability to work out what to do next at an intuitive level, leaving his conscious mind free to focus on matters more difficult and uncertain such as the longer-term tactics and strategy of the race (e.g. how to get ahead of the next car, or judging the best time to pit stop). Indeed, recent research has demonstrated that expert board game players employ just such a division of labour, not just conceptually but physically, utilising the older basal ganglia structure of the brain to intuitively establish their best next move, while using the newer cerebral cortex structure to plan higher level strategy (Wan et al. 2011).

In contrast, at the beginning of their project Watercare's planners were faced with a decision that was both minimally framed and highly complex. They may have been highly experienced and knowledgeable actors, but the situation was nevertheless both uncertain and difficult, thus requiring much explicitation and conscious reflection. In this regard, the two cases were differentiated not just by the *perceived degree of uncertainty* in each domain, but also by *what was uncertain* for the actors in question. For Watercare's planners pretty much everything was uncertain – objectives, evaluation criteria, options, outcomes, etc. – whereas for the race car driver it was the accuracy and precision of his performance which was uncertain. The two cases thus represent different degrees of closure or settlement of decision frames. This is heuristically illustrated in Figure 6.2 by mapping the two cases against the stages of the classical decision process.

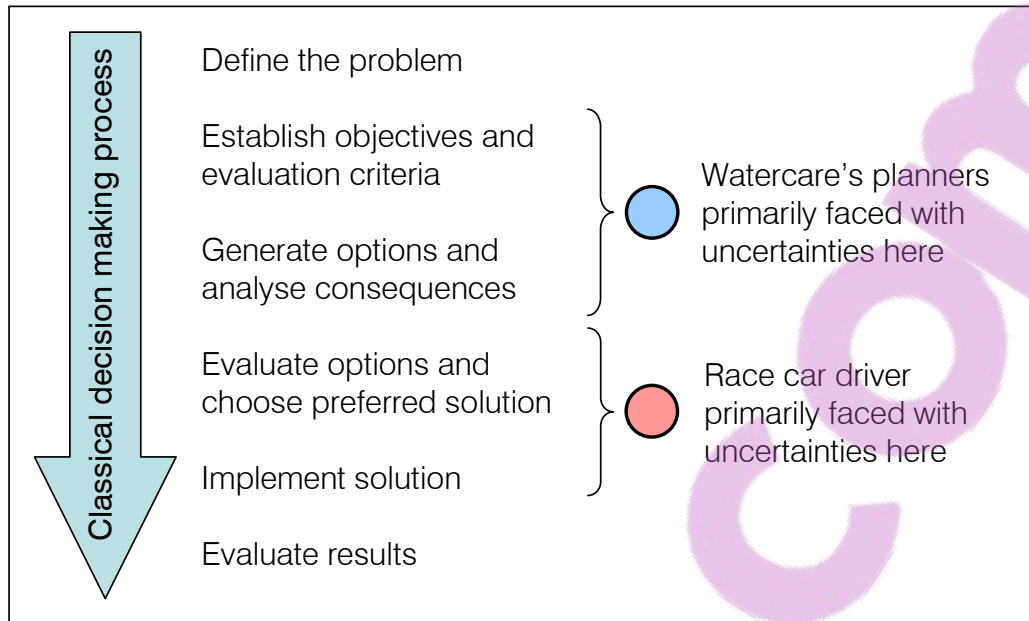


Figure 6.2. Different uncertainties experienced by Watercare's planners and the race car driver

## A typology of problem/decision situations

On this basis then, problem/decision situations may be differentiated according their observed characteristics in three dimensions: (i) the degree to which the problem-solver/decision-maker perceives the problem domain to be uncertain and novel, (ii) the degree of difficulty experienced by the agent with respect to framing and calculating in that domain, and (iii) the nature of the primary uncertainties in that domain. Three general types of situations may be characterised and related in terms of these dimensions (see Table 6.2, page 195, and Figure 6.3, page 196).

**Operational Problems** are situations characterised as routine or familiar to the actor, and of low uncertainty. The actor has significant prior experience of both the specific situation and the calculations to be performed (and in this regard could be referred to as an expert), has all or most of the necessary information about the decision elements, and has access to the resources necessary to proceed. The calculations that the actor must perform may be sufficiently complex that they cannot be performed implicitly, but the actor both knows how to perform the calculations and has access to the calculative resources necessary to carry them out. To the extent that there is uncertainty for the actor, it concerns the exactness (i.e. accuracy and precision) of measurements and calculations. I use the label of "Operational Problem" to highlight the close association with action – in this context the actor either already knows what to do and how to do it, or can rapidly decide the situation with little effort. The earlier case of the race car driver would fall into this category, at the

extreme end of low uncertainty and difficulty (for that driver). Although the watermain planning process would generally be classified as an example of a Planning problem (see below), some aspects that case would also fall into this category, reflecting the planners' skills and experience with certain familiar elements of the process.

**Planning Problems** are situations characterised by some degree of novelty and uncertainty for the actor(s) involved. The actor has some prior experience of similar situations or calculations, but the present context is different in some significant sense. At the outset, the actor is uncertain about some of the decision elements (i.e. the decision frame is significantly incomplete), and about how to proceed in a calculative sense. In this context the actor knows how to go about solving the problem in a general sense (i.e. has a strategy for addressing the situation), but the development of a context-specific methodology is not necessarily straight-forward. The actor must expend a reasonable amount of effort to frame options and evaluate a solution. I refer to these situations as "Planning Problems", since they are the kinds of situations in which actors must explicitly make plans about how to proceed.

**Strategic Problems** are characterised as involving extreme uncertainty, and novelty for the actors in question. Here actors have little or no prior experience with even the general type of situation and the calculations involved. At the outset the problem/decision frame is almost wholly incomplete and there is substantial uncertainty about all elements of the problem. The problem domain is also likely to be perceived as highly complex. In this context the actors may not even know what the problem is, let alone begin to solve it. The problem framing process is difficult in the extreme, characterised by contest and conflict over judgements of the relevance and significance of decision elements, or even over their identity. Such problems may be intractable and the actors may not be able to agree on a frame. Problems of this sort often arise as planning and policy dilemmas with respect to large scale social-environmental problems, and have been referred to in such contexts as "wicked" (Buenaño 1999; Rittel and Webber 1973; Roberts 2001; van Bueren, Klijn, and Koppenjan 2003). In the terms of this discussion, however, I prefer the label of "Strategic Problem" to highlight that it is the problem definition and decision objectives which are in question, first and foremost, as opposed to the solution methodology (Planning problems) or the exactness of information (Operational problems).

The inclusion of the third differentiating dimension (i.e. what is uncertain in the problem domain) points up the possibility that an agent might be reasonably confident about the framing of certain decision elements (e.g. objectives), but remain confused and uncertain about others (e.g. methodology or option details). In this sense, Figure 6.5 does not just define a typology of problem/decision situations, but also the progression of the decision making



process over time. That is, a movement from top-right (high-high) to bottom-left (low-low) on Figure 6.3 illustrates how the agent's overall uncertainty about the problem or decision frame as a whole is resolved through the progression of inquiry and framing (represented by the dark block arrow). Within that broader state of uncertainty, however, the agent may be more or less uncertain about specific decision elements. The resolution of this element-specific uncertainty is heuristically represented by the smaller white block arrows in Figure 6.3, representing nested loops of inquiry and framing within the overall problem-solving process. In this way, strategic problems may be resolved to planning problems, and planning problems to operational problems (although this does not mean that all decisions take this path, or that decision and action are only possible when knowledge is certain).

**Table 6.2.** A basic typology of problem/decision situations (see also Figure 6.5)

	<b>Operational Problems</b>	<b>Planning Problems</b>	<b>Strategic Problems</b>
Characteristics of the problem domain	Relatively low uncertainty and novelty. Agent has all necessary experience, information, knowledge, and resources.	Moderate uncertainty and novelty. Agent has generally relevant experience, and access to some information, knowledge, and resources.	Extreme uncertainty and novelty. Agent is a novice and is significantly constrained by lack of information, knowledge, and resources.
Characteristics of agent's performance (framing and calculating)	Agent has little difficulty with framing and calculating. Moves rapidly to action.	Agent has some difficulty with framing and calculating. Action is preceded by reflective as well as merely calculative processes.	Agent has extreme difficulty with framing and calculating. Process is likely to be messy, confusing, and protracted. Agent must muddle.
Nature of uncertainties	Technical: agent knows what must be achieved and how to achieve it. Uncertainties concern exactness (accuracy and precision) of measurements and calculations.	Methodological: Agent knows what must be achieved but not specifically how. Uncertainties concern the interpretation of objectives and the appropriateness of solution methodologies.	Definitional: Agent is uncertain about decision objectives; may even be uncertain about the existence and nature of the problem.

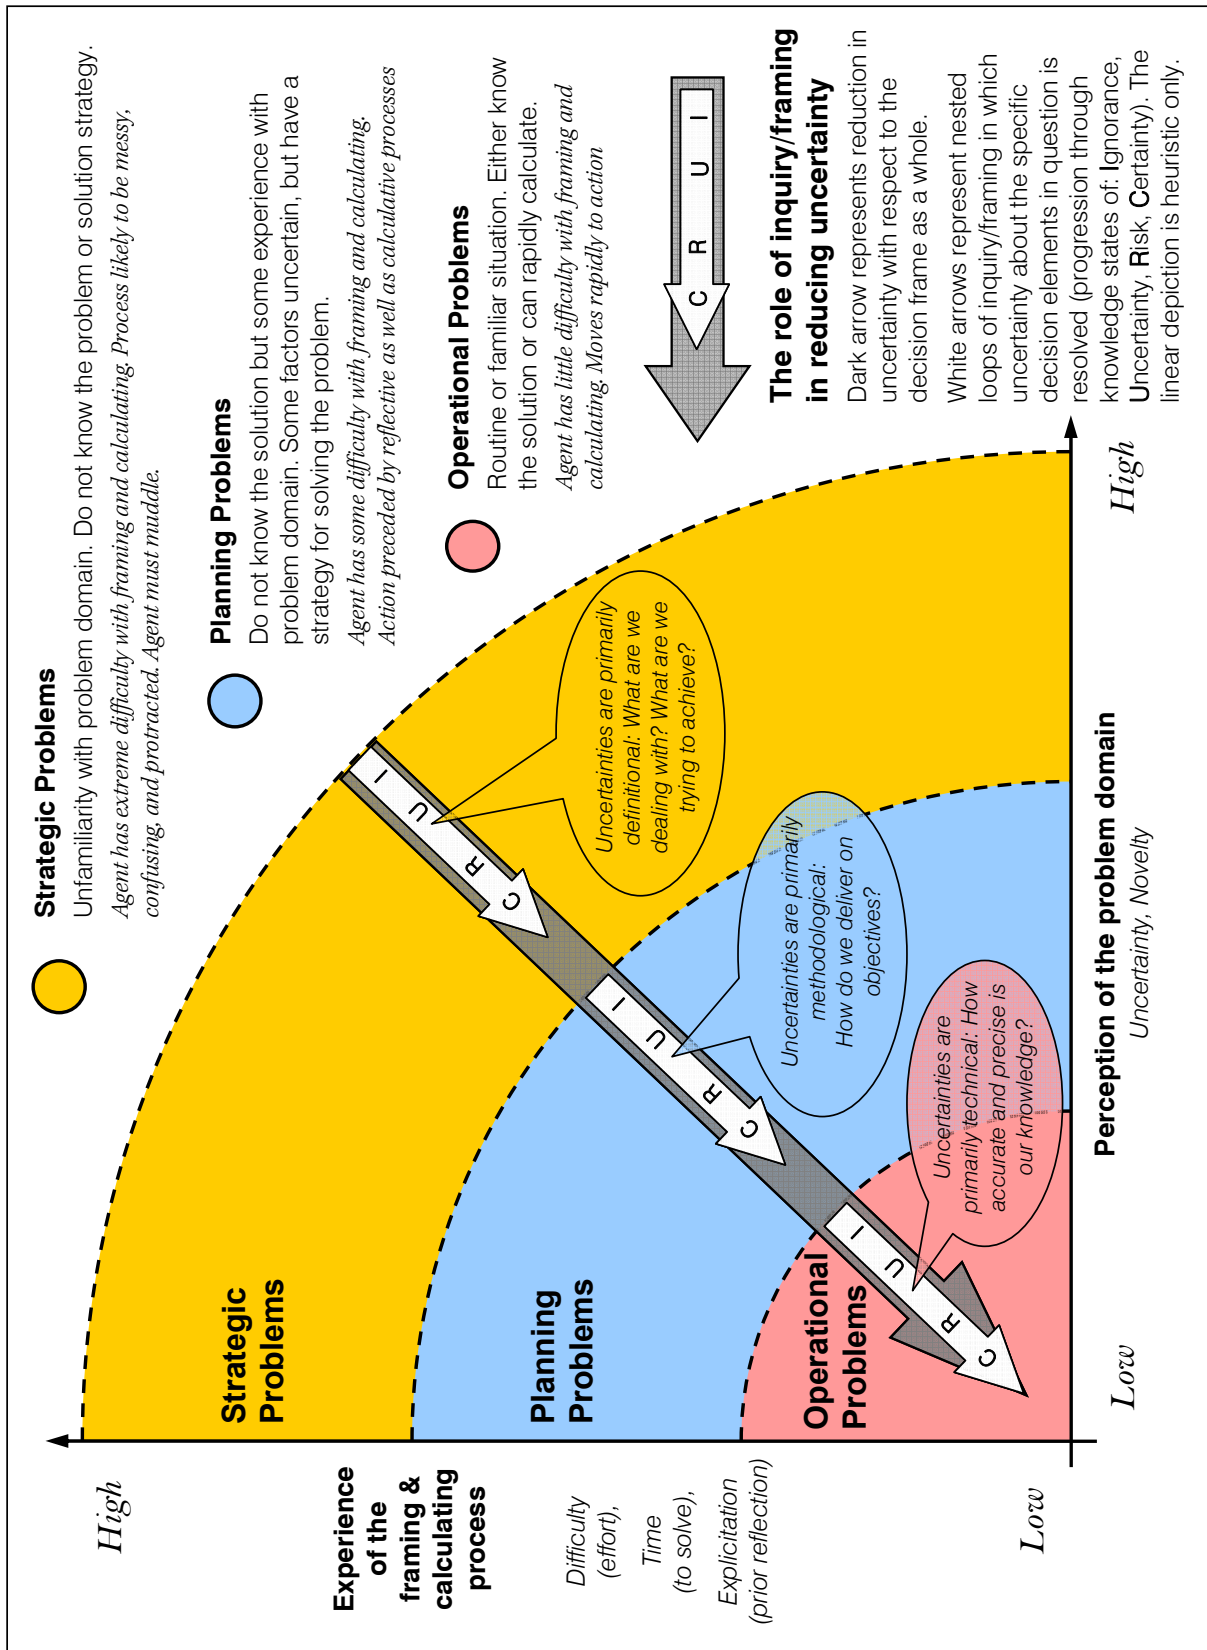


Figure 6.3. Problem situations defined in terms of the agent-context relationship.

The x-axis represents the problem situation, as it is perceived by the agent, and the y-axis the nature of the problem framing / solving process, as it is experienced by the agent. The axes, and the problem definitions, are relative (i.e. subjective) to the agent.

## Extending the typology

In 1993, Funtowicz and Ravetz published a now well cited paper in which they conceptualised and differentiated modes of scientific inquiry. The authors were motivated by a concern that traditional (normal) science could not adequately cope with the uncertainties and ethical dilemmas presented by socio-environmental problems (Funtowicz and Ravetz 1993). They argued that the policy issues presented by such problems had common features which distinguished them from traditional scientific problems, including the inability of normal science to provide factual predictions, the magnitude, scale, and ethical complexity of outcomes, and the urgency of policy decisions (Funtowicz and Ravetz 1993, p 742). In the face of such complex, large-scale problems, coping with significant systems uncertainties and assuring the quality of scientific information were emerging as the central integrating concepts of a new post-normal scientific methodology (Funtowicz and Ravetz 1993, p 742).

In order to conceptualise that methodology and differentiate it from more traditional modes of scientific problem solving, Funtowicz and Ravetz (1993, p 744) characterised a series of problem situations according to the degree of intensity along two dimensions: *Systems uncertainties*, which described the epistemological and ethical uncertainties which arise in regard to the comprehension and management of inherently complex social and environmental systems; and *Decision stakes*, which described the conflicting purposes and value commitments involved through multiple stakeholders and interests. They noted in this regard that their framework showed “the interaction of the epistemic (knowledge) and axiological (values) aspects of scientific problems... We notice that uncertainty and decision stakes are the opposites of attributes which had traditionally been thought to characterize science, namely its certainty, and its value neutrality” (Funtowicz and Ravetz 1993, p 744). They labelled the three problem situations for the corresponding science methodologies (applied science, professional consultancy, post-normal science) which they believed were appropriate for resolving those sorts of problems.

The typology of problem/decision situations represented in Table 6.2 and Figure 6.5 bears resemblance to the framework published by Funtowicz and Ravetz. They both characterise and differentiate problem/decision situations in terms relative to the problem solver(s), and both employ similar tripartite classifications based on the common denominator of uncertainty. There are three additional features of latter framework that I wish to capture here in order to extend and enrich the characterisation of problem/decision contexts presented above. The first is to explicitly acknowledge that the framing of problems and decisions (and not just those of a scientific nature) is both an epistemic (knowledge-based)

and axiological (values-based) task. That is framing involves identifying and knowing decision elements, in an objective, scientific sense, which involves uncertainties of an epistemological nature; roughly corresponding to Funtowicz and Ravetz's notion of Systems Uncertainties. Framing also involves judging the relevance and significance of elements to the problem or decision at hand (this is explained further later), which involves uncertainties of an ethical and value-laden nature; roughly corresponding to Funtowicz and Ravetz's notion of Decision Stakes. In this regard, the definitions of the three general problem types (Operational, Planning, Strategic) can be extended to reflect the nature of the characteristic uncertainties in these dimensions (see Table 6.3).

**Table 6.3.** Nature of epistemological and axiological uncertainties defining problem domains  
(Extends Table 6.2 by reference to Funtowicz and Ravetz 1993)

	<b>Epistemological Uncertainty</b>	<b>Axiological Uncertainty</b>
	<i>Represents the technical, methodological, and ontological uncertainties which arise in regard to the comprehension of inherently complex realities.</i>	<i>Represents the evaluative and strategic uncertainties (value and ethical dilemmas) which arise in resolving the conflicting interests of multiple stakeholders and the unequal distribution of costs and benefits.</i>
<b>Operational Problems</b>	<u>Technical</u> : uncertainties concern the inexactness (i.e. accuracy and precision) of performance, measurement, and calculation.	None (the problem, purpose, and solution methodology are not in question)
<b>Planning Problems</b>	<u>Methodological</u> : uncertainties concern the reliability of information and theories.	<u>Interpretive</u> : uncertainties concern the interpretation of objectives (by agents) and the appropriateness of solution methodologies.
<b>Strategic Problems</b>	<u>Ontological</u> : uncertainties concern the appropriateness and legitimacy of competing disciplinary and lay viewpoints for rendering the world knowable.	<u>Strategic</u> : uncertainties concern the existence and nature of the problem (resolution of conflicting problem perceptions), and the specification and agreement of decision objectives (resolution of conflicting interests).

However, while it is useful to acknowledge the epistemological and axiological aspects of problem/decision framing, it would be inaccurate to represent them as separate and independent. For instance, I am not sure that what I have called Operational problems involve more than trivial uncertainties in the axiological dimension at all. Rather, the kinds of value dilemmas and ethical questions which the axiological dimension is intended to represent do not arise until outcomes, methods, and objectives are called into question. And, further, as uncertainty about outcomes, methods, and objectives increases, it may not be possible to strictly separate the epistemological and axiological aspects of that uncertainty.

That is, where uncertainty is extreme, knowing entities may be intricately tied up with judging (valuing) them, such that the traditional distinction between facts and values cannot realistically be maintained (Funtowicz and Ravetz 1993; Megill 1994).

The second feature that I wish to capture is the association between problem type and the increasing number and diversity of stakeholder groups with interests in the problem solving task. Funtowicz and Ravetz (1993) discussed that as decision stakes increase so too does the number and diversity of stakeholders both seeking to define the purposes of the decision (which will increasingly be in conflict), and affected by its outcomes. Similarly, increasing systems uncertainties are associated with an increasingly diverse peer community seeking to evaluate the quality of the problem solving task. This notion of an expanding stakeholder community emphasises the increasing significance of the axiological dimension as a source of uncertainty within the problem domain; reflected in the growing conflict between the purposes and interests of disparate stakeholders, and in the increasingly general focus of quality assurance due to the lack of relevant technical expertise within the extended peer community.

The third feature that I wish to highlight is how the three problem types stand in relation to each other, in terms of (i) the respective roles fulfilled by the problem solver(s) and stakeholder(s), (ii) the translation of interests that occurs as problems are resolved from Strategic to Operational levels, and (iii) the transmission of information between problem levels. For each problem category, Funtowicz and Ravetz distinguished the internal problem solver role from the external role of the stakeholder(s) who defines purposes and evaluates the quality of the problem solving task (see Table 6.4).

*Table 6.4. System roles in Funtowicz and Ravetz (1993)*

<b>Problem category</b>	<b>Internal Role</b> <i>Problem solver / Decision maker</i>	<b>External Roles</b> <i>Defining purpose, evaluating quality</i>
Applied science	Researcher	Manager, Users
Professional consultancy	Professional consultant	Client, External (affected) stakeholders
Post-normal science	Policy maker	Extended peer community (participant stakeholders)

The hierarchical differentiation of problem solver and stakeholder roles in Funtowicz and Ravetz (1993) is a common feature of various system-theoretic representations of organisational and societal systems (e.g. Beer's Viable Systems Model of organisations, and

evolutionary models of Large Technical Systems; see Appendix II). In the normative perspective Sheffield and Guo (2007) synthesised the ideas of three prominent twentieth century philosophers to structure a hierarchical representation of knowledge management and organisational systems. Each level in their hierarchy was defined in terms of a problem solving / decision making role (Churchman's system roles), a knowledge perspective which informs that role (Habermas' knowledge interests), and a type of work which is performed by that role (Ulrich's boundary judgements). The three internal system roles were (Sheffield and Gou 2007, p 76):

- The client: whose interests and values are to be served by the system, and who contributes the necessary purpose and direction, including the standards by which the performance of the system will be judged. The client is associated with the personal knowledge perspective; the objectives of the system are framed from the client's personal value system and the confirmation (or disconfirmation) of the achievement of the system objectives is subjective to the client's perception.
- The decision maker: who, as a source of control/power, contributes the necessary means, resources, and decision authority to the system. The decision maker is associated with the organisational knowledge perspective; he or she controls the resources of the system and makes decisions as to how those resources should be employed to give effect to the client's objectives and values.
- The planner: who is the source of technical expertise within the system, contributes the knowledge and skills to perform the technical, productive work of the system. The planner understands the intentions of the decision maker and has the necessary factual knowledge and technical expertise to organise and perform empirical inquiry and implementation. The planner also measures technical performance and provides feedback to the decision maker.

Sheffield and Guo (2007, p 76) also defined a fourth role, external to system, which they labelled "the Witness". This role was not directly involved in problem solving or decision making activity within the system, but was nevertheless affected by the outcomes of such activity. It was therefore defined in terms of a critical or ethical consideration, and as either contributing to or countering the legitimacy of the system (Sheffield and Gou 2007). In this sense, the role of "the Witness" parallels Funtowicz and Ravetz's (1993) "external stakeholder" and "extended peer community" roles.

A shared feature of the descriptions of the problem solving and stakeholder roles at each

problem level, in both Funtowicz and Ravetz (1993), and Sheffield and Guo (2007), is the dual functions of those roles in terms of the translation of interests and transmission of information between and through those roles. These functions, which I call “directive” and “evaluative”, respectively, are also evident in other systems representations of organisational and societal systems (e.g. Beer’s Viable Systems Model of organisations, and evolutionary models of Large Technical Systems; see Appendix II). Table 6.5 elaborates the functions with respect to each role (Table 6.5 also renames the roles at each level to better suit the ultimate purpose of this framework)

**Table 6.5.** Directive and Evaluative functions of problem-solver and stakeholder roles

<b>Problem domain</b>	<b>Role</b>	<b>Directive function</b> <i>Translation of interests</i>	<b>Evaluative function</b> <i>Transmission of information</i>
External	External (affected) Stakeholder	Translates interests into boundary criteria for the legitimate (ethical) behaviour the system.	Evaluates the ethical behaviour (legitimacy) of the system.
Strategic	Client / Owner	Translates his or her interests into the purposes of the system (definition of objectives, and performance criteria).	Evaluates the Strategic performance of the system and reports or publicises information about that performance outside the system.
Planning	Manager / Planner	Interprets the Client’s objectives within an organisational context and translates them into specific plans to be actioned, and parameters for acceptable performance.	Evaluates the Operational performance of the system against objectives and reports to the Client / Owner on plans actioned and the relative fulfilment of objectives.
Operational	Operator / Primary Producer	Translates the Manager / Planner’s plans into specific actions within a specific functional and technical context, and performs those actions.	Measures the technical performance of the system and reports to the Manager / Planner on actions taken and performance outcomes.

Sources for Table 6.5: (Funtowicz and Ravetz 1993; Sheffield and Gou 2007)

### **An extended typology of problem domains**

Table 6.6, overpage, combines the problem categories (Operational, Planning, Strategic; Table 6.2), the problem-solver and stakeholder roles (Operator/Primary Producer, Manager/Planner, Client/Owner; Table 6.5), and the functions of those roles (Directive and Evaluative; Table 6.5) to create an extended typology of the subjective relationship between the agent and the problem solving/decision making context.

This page intentionally blank



**Table 6.6.** Framework characterising the general features of *problem* domains, the *problem solver* roles associated within those domains, and the functions of those roles

<b>Problem type:</b>		<b>Operational</b>	<b>Planning</b>	<b>Strategic</b>	<b>Notes</b>
Key Characteristics	<i>Of the problem domain</i>	Relatively low uncertainty and novelty. Agent has all necessary experience, information, knowledge, and resources.	Moderate uncertainty and novelty. Agent has generally relevant experience, and access to some information, knowledge, and resources.	Extreme uncertainty and novelty. Agent is a novice and is significantly constrained by lack of information, knowledge, and resources.	<i>Represents how an actor perceives the situation in terms of the degree of uncertainty (extent to which decision elements are unknown or unknowable by the actor), and novelty (extent to which the actor is unfamiliar with the situation and the calculations he or she has to perform).</i>
	<i>Of the agent's problem solving performance</i>	Agent has little difficulty with framing and calculating. Moves rapidly to action.	Agent has some difficulty with framing and calculating. Action is preceded by reflective as well as merely calculative processes.	Agent has extreme difficulty with framing and calculating. Process is likely to be messy, confusing, and protracted. Agent must muddle.	<i>Represents how the actor experiences the problem-solving/decision-making process, in terms of the difficulty of that process, the duration, and the degree of explicitation and prior reflection involved.</i>
Key Uncertainties	<i>Epistemological (knowledge) dimension</i>	<u>Technical</u> : uncertainties concern the inexactness (i.e. accuracy and precision) of performance, measurements, and calculations.	<u>Methodological</u> : uncertainties concern the reliability of information, theories, and methodologies.	<u>Ontological</u> : uncertainties concern the appropriateness and legitimacy of competing disciplinary and lay viewpoints for rendering the world knowable.	<i>Represents the technical, methodological, and ontological uncertainties which arise in regard to the comprehension of inherently complex realities.</i>
	<i>Axiological (values) dimension</i>	None (the problem, purpose, and solution methodology are not in question)	<u>Interpretive</u> : uncertainties concern the interpretation of objectives (by agents) and the appropriateness of solution methodologies.	<u>Strategic</u> : uncertainties concern the existence and nature of the problem (resolution of conflicting problem perceptions), and the specification and agreement of decision objectives (resolution of conflicting interests).	<i>Represents the evaluative and strategic uncertainties (value and ethical dilemmas) which arise in resolving the conflicting interests of multiple stakeholders and the unequal distribution of costs and benefits.</i>
<b>Problem solver role:</b>		<b>Technical expert / Primary producer</b>	<b>Manager / Planner</b>	<b>Client / Owner</b>	<b>Notes</b>
Directive function of the problem solver role	<i>Directive Function</i>	Translates the Manager / Planner's plans into specific actions within a specific functional and technical context, and performs those actions.	Interprets the Client's objectives within an organisational context and translates them into specific plans to be actioned, and parameters for acceptable performance.	Translates his or her interests into the purposes of the system (definition of objectives, and performance criteria).	<i>Represents the translation of interests between and through roles. Higher level roles define the purposes and boundary conditions of lower level roles.</i>
	<i>Questions defining the directive function</i>	What actions need to be taken to maintain or correct performance?	How should the objectives be interpreted in specific contexts? What do they mean in practice?  How can the objectives be reliably achieved? What methodologies are appropriate in specific contexts?	What is the current state of the world/system? Why is it problematic?  What is the desired future state of the world/system? Why is it desirable?  What are the objectives? Why should they be pursued?  What constitutes achievement of the objectives?	<u>External Stakeholder Role</u>  <i>An additional External (affected) Stakeholder role may be conceptualised as external to the problem-solving system, but affected by the actions of that system.</i>  <i>The External Stakeholder may seek to translate its interests into boundary criteria for the legitimate (ethical) behaviour of the system:</i>  <u>Questions:</u> What constitutes acceptable (good, right) behaviour?  <u>Outputs:</u> Conditions and parameters of acceptable behaviour
	<i>Outputs (arrows indicate direction of conduct)</i>	Actions performed	Action plans developed, performance parameters specified (boundary conditions for action)	Objectives and criteria for interpretation specified (boundary conditions for methodology)	
Evaluative function of the problem solver role	<i>Evaluative function</i>	Measures the technical performance of the system and reports to the Manager / Planner on actions taken and performance outcomes.	Evaluates the Operational performance of the system against objectives and reports to the Client / Owner on plans actioned and the relative fulfilment of objectives.	Evaluates the Strategic performance of the system and reports or publicises information about that performance outside the system.	<i>Represents the transmission of information between roles, and the evaluation of that information within roles. Each role records and reports on its performance to higher level roles which evaluate that performance.</i>
	<i>Questions defining the evaluative function</i>	How is system performing relative to performance parameters?  How accurate and precise is our information?	Are objectives being achieved? If not, why not?  How reliable is our knowledge of the world/system?	Have objectives been achieved?  Did they result in the desired outcomes?	<u>External Stakeholder Role</u>  <i>The External Stakeholder may seek to evaluate the performance of the problem-solving system in terms of ethical/critical considerations, i.e. How good or bad was the performance of the system?</i>
	<i>Outputs (arrows indicate direction of conduct)</i>	Information on actions taken and performance outcomes	Information on plans actioned and achievement of objectives	Information on performance of the system	

This page intentionally blank

## Reconceptualising “decision” in performative terms

It is clear from the preceding discussion that real decision making behaviour is only approximate in form to the traditional ideal of the rational choice in a limited set of circumstances; i.e. as an explicit procedure by which an agent identifies some preferred course of action which optimally satisfies all decision criteria and rules. This implies that it is necessary to disconnect the normative criteria of rationality from the assumed form of a decision. To achieve this I borrow Callon’s (1998a, 1998b, 1999) notion of “framing” as a process of disentangling, from everything else, what is relevant and significant for the task at hand, to explore the question of what it means to settle a decision frame under conditions of the rational imperative (i.e. that one must always do what is best). That is, if frames must be constructed as a series of distinctions about what is relevant and significant, then what are the conditions for settling decision frames rationally? The discussion does not try to define abstract parameters or conditions of rationality against which real decision making might be judged, since such a task would, I think, be a case of *reductio ad absurdum*. Rather, the analysis calls attention to those sources of uncertainty which are common to rational calculations everywhere, since helping actors to resolve judgements with respect to those uncertainties would be the goal of decision support interventions. The reader is referred to Appendix VIII which presents a more detailed conceptual development of the conditions for cutting off the framing process in a reasonable manner.

### Framing as constitutive of decision

In its broadest normative sense, the rational imperative requires that agents should do what is appropriate with respect to their beliefs and actions (Rescher 1995). In this sense, rationality (reason) refers to a system of norms for the selection of beliefs and actions; i.e. it is concerned with what one ought to believe and what one ought to do. Traditionally, the distinction has been made as follows, between the rationality of beliefs and the rationality of actions:

Theoretical reason... addresses the considerations that recommend accepting particular claims as to what is or is not the case. That is, it involves reflection with an eye to the truth of propositions, and the reasons for belief in which it deals are considerations that speak in favor of such propositions' being true, or worthy of acceptance. Practical reason, by contrast, is concerned not with the truth of propositions but with the desirability or value of actions. The reasons in which it deals are considerations that speak in favor of particular actions being good, or worthy of performance in some way. (excerpt from Wallace 2009)

The distinction points up that decision making involves considerations in both epistemological and axiological dimensions (Funtowicz and Ravetz 1993). That is, in making decisions agents may consider the reliability of various kinds of evidence and claims as to the nature of the world around them (which results in changes in the agents' beliefs), while also weighing the relative significance (value) of those things that they take into account (which results in changes in the agents' intentions with respect to future actions). It also highlights that in order to form rational beliefs and intentions agents must calculate (Rescher 1995; Callon 1998a). That is, where alternative possibilities for belief and/or action exist, and are perceived by the agent in question, then the rational choice of one possibility over the others requires the agent to reject those other possibilities as less appropriate. It is this notion of a prior intelligent calculation distinguishing a particular belief or intention from amongst various possibilities which separates rational behaviour from mere impulse or whim.

The formation of rational beliefs or intentions involves two types of calculation, corresponding to the two categories of reasoning, above. In the case of theoretical reasoning, an agent may be justified in believing that a proposition,  $p$ , is true if the agent has adequate indication, in the absence of overriding evidence to the contrary, that  $p$  is true ('epistemology' in the Cambridge Dictionary of Philosophy 1999; refer also to Appendix VIII). Thus forming a rational belief requires the agent to weigh the relative strength (i.e. reliability) of the evidence for and against that belief. In the case of practical reasoning, an agent may be justified in his or her intent to act if the intended action is expected to be maximally efficient with respect to achieving the agent's objectives. Thus forming rational intent requires the agent to weigh the relative significance (value) of outcomes of alternative courses of action (accounting for outcomes desired, outcomes undesired, and outcomes not undesired). The latter calculation is clearly predicated on the former, since one cannot reasonably attach a value to something if one does not believe it to be the case.

These calculations are only possible when the objects of the calculations (i.e. those things which are to be taken into account) have been identified, specified, valued, and related to those other things which are also to be counted, as well as to those which will be left out of the calculation (Callon 1998a, 1998b, 1999; Callon and Muniesa 2005; Latour 1987). That is, one cannot weigh evidence unless one has the evidence at hand, and one cannot weigh outcomes unless one knows what those outcomes will be and their relative significance. Thus, in order to perform calculations, agents must engage in a certain amount of prior work, both cognitive and material, to identify and sort out those things which are relevant and significant for the calculation at hand, from those which are not (Callon 1998b, 1998a, 1999; Callon and Muniesa 2005). This work has been called "framing" (i.e. the construction of

problem or decision frames; Buenaño 1999; Callon 1998b, 1998a, 1999; Chia 1994; Cooper 1986; Herbst 1976).

Callon (1998a) argued that the framing of any thing (entity) which is to be taken into account in the formation of beliefs or intents necessarily involves three distinctions. First, whatever is to be taken into account must be recognised, its existence must be both perceived and proven. In other words, it must be known. This requires that the entity be observable by someone and, in most cases, also describable, but does not mean that the entity must be a physical object. Less tangible entities, such as objectives and preferences, are no more or less difficult to observe and describe than physical objects. Where conflicting evidence or competing knowledge claims exist, then these must be evaluated to establish which is more reliable. Second, whatever is to be taken into account must be judged as significant in some way by at least one of the decision stakeholders. This requires that both the entity (i.e., its magnitude) and its significance (i.e., its value; which may not necessarily be proportional to its magnitude) must be made measurable and comparable with other entities. Where competing judgements exist then these must be resolved so that valuations are consistent. Third, the relevance of things which are to be taken into account must be established. This means that each entity must be somehow related to or associated with those other things which are to be taken into account, and arranged so that calculative operations may be performed. Relevance may be established on the basis of physical and/or value association, while the arrangement of things to be calculated may be performed physically (such as on an invoice or a trading screen, or in a shopping cart), and/or cognitively, in both conscious and sub-conscious senses (Callon 1998a, 1998b, 1999; Callon and Muniesa 2005, pp 1230-1232). In this sense, framing may be understood as constitutive of knowledge; i.e. the making of cognitive and material distinctions in a particular context, which order the world into “a version of reality to which we subsequently respond” (Bell 1999; Chia 1994, p 781; Tsoukas and Vladimirou 2001; see also the discussion of the pragmatist conception of knowledge in Chapter 2).

## **Sources of uncertainty in decision frames**

In order to perform this framing work agents must employ or mobilise a variety of boundary-objects (Bowker and Star 1999; Callon 1999; Latour 1987). For instance, language, often in particular forms, is necessary to describe entities; specific tools, methods, and methodologies will be necessary to measure and quantify them; while in some contexts political agreement between agents may be necessary to recognise an entity as significant in some form or other.

This means that quite specific work must be performed to construct and extend decision frames, involving the expenditure of effort and the mobilisation of resources, both material and political (Callon 1998a; Callon and Muniesa 2005; Latour 1987). But while boundary objects make possible the stabilisation and framing of entities, they also simultaneously provide “an opening onto other worlds, thus constituting leakage points where overflowing can occur” (Callon 1999, p 188). Every decision frame is, in this regard, subject to overflows, or what are more commonly referred to as externalities:

What economists say when they study externalities is precisely that this work of cleansing, of disconnection, in short, of framing, is never over and that in reality it is impossible to take it to a conclusion. There are always relations which defy framing. It is for these relations which remain outside the frame that economists reserve the term externalities. The latter denotes everything which the agents do not take into account and which enables them to conclude their calculations. But one needs to go further than that. When, after having identified them, the agents... decide to reframe them – in other words to internalize the externalities – other externalities appear. I would suggest the term ‘overflowing’ to denote this impossibility of total framing. (*excerpt from Callon 1999, p 188*)

This overflowing occurs due to the limitations of the ways in which humans come to know and judge the world (Loasby 1999; Rescher 1995, 2005), and the impossibility of ever being able to absolutely prioritise everything for the purposes of rational calculation (Callon 1998b; Lindblom 1959, 1979; Loasby 1999). If overflows are always present, then all problem/decision frames are necessarily subject to a degree of uncertainty, no matter how small. This is because the relations or associations which overflow the frame are also conduits through which translations may take place. They provide openings through which external mediators can influence or alter the entities within the frame, thus constituting possibilities for framed entities to become other than what they have been taken to be (see Chapter 2 for the meaning of “translations” and “mediators” in ANT terminology). Overflows therefore constitute sources of uncertainty with respect to the decision frame itself (where uncertainty in this sense is epistemic not aleatory; Hora 1996).

Conceptually, each of the distinctions which constitute the decision frame embodies possibilities for things to be other than what they are taken to be. The distinction of recognition embodies the potential for ambivalence about the objective nature of things in the world, and about what will happen in the future; arising from the extent and reliability of information and knowledge about the problem domain, and about the calculations that need to be made in that domain (including experiential knowledge). The distinction of significance embodies the potential for equivocality about the value of things in the world; arising from the effects of bias, caprice, or conflict in value judgements. The distinction of relevance

embodies the potential for unstable or mutable relations; arising from errors or inconsistencies in the way things are arranged for the purposes of calculation. Finally, since human reasoning is not infallible, the agent's own information processing capabilities embody the potential for error (i.e. is the agent a reliable calculator or is the agent prone to making various sorts of errors such as bias, assumptions, mistakes?). These sources of uncertainty in decision frames are summarised in Table 6.7.

In this regard, framing may be understood as the work that agents perform to identify and frame overflows in order to reduce uncertainty about beliefs and intentions. But if decision frames can never be absolute then this implies that there may be situations in which the agent cannot complete the framing process with any certainty, and may have to proceed on the basis of incomplete knowledge (uncertainty) or even under ignorance. Rationality requires that when agents calculate under these conditions they should account for the implications of the possibilities of things being other than what they are taken to be. Traditional decision theory prescribes various rules for how agents should proceed in these circumstances (see for instance Hansson 1994 for an overview; refer also to Appendix VIII).

**Table 6.7.** Types and sources of epistemic uncertainty in decision frames

<b>Framing distinctions</b>	<b>Type and source of uncertainty</b>
Recognition	<i>Ambivalence about the objective nature of things in the world, and about what will happen in the future.</i> This type of uncertainty arises from the extent and reliability of information and knowledge about the problem domain, and about the calculations that need to be made in that domain (including experiential knowledge).
Significance	<i>Equivocality about the value of things in the world.</i> This type of uncertainty arises due to bias, caprice, or conflict in the formation of value judgements
Relevance	<i>Mutability of relations between things.</i> This type of uncertainty arises due to errors or inconsistencies in the way things are arranged in relation to each other for the purposes of calculation.
All	<i>Fallibility of cognition.</i> This type of uncertainty arises from the reliability of agents' information processing capabilities (i.e. is the agent a reliable calculator or is the agent prone to making various sorts of errors such as bias, assumptions, mistakes?)

## Conditions for cutting off framing reasonably

If decision frames can never be absolute then this implies that agents must cut off the framing process at some point so that they may proceed with their cognitive and practical endeavours; else it would never be possible for agents to establish belief or intent (this is the concept of satisficing more broadly stated; Simon 1955, 1957). In principle this means that,

unless the limits of time and resources are confronted, agents should only terminate framing when they rationally believe that further framing will be unlikely to alter their beliefs or intentions in a significant way. This implies that agents should constantly evaluate whether continuing the framing process will be likely to yield a decision significantly different from that which is currently framed, taking into account the future costs of framing. The agent's expectations in this regard will depend on the extent and duration of the agent's search to date (i.e. what and how much he already knows about the broader problem domain), his experience with the specific framing process so far, and with the type of problem being addressed (as a guide to how easy or difficult it might be continue), whether he has the time and resources to continue, and the nature and reliability of the agent's process of reasoning from this evidence to an expectation about the potential success of continuing the framing process. (The conditions for cutting off the framing process in a reasonable manner are developed in Appendix VIII).

## **Decision support strategies for Chief Risk Officers**

Reason requires that agents should do what is appropriate with respect to belief and action. Where alternative possibilities for belief and/or action exist, and are perceived, the rational choice of one possibility over the others requires a calculation as to which is most appropriate. In the theoretical sense, reason requires that agents should weigh the relative strength (i.e. weight and reliability) of evidence for and against the truth of things in order to establish rational belief as to what is or is not the case. In the practical sense, reason requires that agents should weigh the relative significance (i.e. value) of things in order to establish a rational intent to act. This latter calculation is necessarily predicated on the former because one cannot reasonably attach a value to something if one does not believe it to be the case. Reason requires that both calculations should be consistent and account for uncertainty where necessary. These normative requirements for rational calculation are summarised in the left-hand columns of Table 6.8.

Calculation is predicated on prior distinctions to identify and sort out what is relevant and significant for the calculation from what is not. The key framing distinctions for any decision are summarised in the middle column of Table 6.8. Framing begins with the mere perception of the world around us but can never be completed in an absolute sense due to the limitations of the ways in which humans come to know and judge the world and the impossibility of ever being able to absolutely prioritise everything for the purposes of rational



calculation. Framing always takes place in a context of certain inputs and other factors which influence agents' abilities to frame and calculate. Some inputs and factors are agent-specific, while others are context-specific. Collectively they influence how an agent may perceive a given situation and that agent's subsequent decision making experience. The various inputs and factors which characterise the agent-context relationship are summarised in the right-hand columns of Table 6.8.

Consequently framing is always more or less difficult, costly, and time consuming, and the distinctions which constitute decision frames are always subject to some degree of uncertainty, however small. In principle this means that, unless the limits of time and resources are confronted, agents should only terminate framing when they rationally believe that further framing will be unlikely to alter their beliefs or intentions in a significant way. This implies that agents should constantly evaluate whether continuing the framing process will yield a better decision than that which is currently framed, taking into account the future costs of framing.

However, it is a significant feature of real-world decision making that beliefs and intentions, and hence the distinctions and calculations which constitute them, only become matters for conscious reflection when they are called into question; that is, when the agent perceives some significant uncertainty or difficulty in making those distinctions or calculations. In the absence of perceived uncertainty or difficulty, humans tend to move automatically from evidence and values to beliefs and intentions. Thus, unless we perceive a reason to doubt it, we tend to trust the information or knowledge at our disposal as being reliable. This presumptive use of trust is pragmatically necessary for efficient action, in both cognitive and practical senses, but is always provisional. It may be revoked if things go awry or we become aware of a reason to be uncertain (see Chapter 2).

The following sections draw on the earlier work in this chapter to conceptualise three generic strategies by which Chief Risk Officers might help agents to make more rational decisions. It is argued that, due to certain unique constraints on the CRO role, CROs should focus their attention, vis-à-vis the decision support function of ERM, on one particular methodological approach (two other approaches are also identified, but are dismissed as unsuitable as *general* decision support strategies for CROs). The methodological approach is not defined in detail. Rather, the discussion focuses on identifying those features which are likely to be common to the approach across different contexts and those factors which are likely to differentiate the approach across contexts. In this regard, the chapter presents a strategic outline of the methodological approach which could serve as a stepping off point for future research and development in this area.

**Table 6.8.** A summary of decision making

Reason requires that agents should do what is appropriate with respect to belief and action. Where alternative possibilities for belief and/or action exist, and are perceived, the rational choice of one possibility over the others requires a calculation as to which is most appropriate. Calculations should be consistent and account for uncertainty where necessary.		Calculation depends on prior distinctions to identify and sort out what is relevant and significant from what is not. Called “framing”, this work is both cognitive and material in nature.	Framing always takes place in a context of certain inputs and other factors which influence agents’ abilities to frame and calculate. Some inputs and factors are agent-specific, while others are context-specific. Collectively they influence how an agent may perceive a given situation and that agent’s subsequent decision making experience.	
Normative requirements	Calculations	Framing distinctions	Inputs	Influencing factors
<p><b>Consistency:</b> beliefs and intentions should cohere. That is, they should adhere to basic rules of logic; should not conflict with knowledge or preferences; and should be consistent (i.e. not arbitrarily changed due to immaterial factors)</p> <p><b>Accounting for uncertainty:</b> in situations of risk, uncertainty, or ignorance, calculations should properly account for the possibility or probability that things might be otherwise.</p>	<p><b>Belief:</b> rational belief is achieved by weighing the relative strength (i.e. weight and reliability) of evidence for and against the truth of things.</p> <p><b>Intent:</b> rational intent is achieved by weighing the relative significance (value) of things.</p>	<p><b>Recognition:</b> recognising the existence, nature, and magnitude of things to be taken into account. Where necessary, judging the reliability of competing evidence or knowledge claims.</p> <p><b>Significance:</b> judging the value of things to be taken into account. Where necessary, resolving competing value claims.</p> <p><b>Relevance:</b> relating those things which are to be taken into account with each other and to those things which will be left out of the calculation. Arranging things so that calculative operations may be performed.</p>	<p>Information about/ knowledge of things and the world or system in which they are constituted.</p> <p>Information about/ knowledge of theory, methods, and procedures relevant to performing calculations.</p> <p>Experience with performing the same or similar calculations in same or similar situations elsewhere.</p> <p>Personal and cultural values, interests, and world views.</p>	<p>Available time and resources (investigative, metrological, calculative)</p> <p>Cognitive abilities of actors</p> <p>Complexity of calculations</p> <p>Motivation of actors</p> <p>Emotional experience and mood state of actors</p> <p><i>And where more than one actor is involved:</i></p> <p>Group cultural and emotional dynamics</p> <p>Efficiency and effectiveness of communication</p> <p>Ability of political actors to reach agreement</p>

## Strategy 1: Formatting the rational decision maker

The first possible strategy is conceived on the rationale that rational decision making “can exist” or at least be approximated in practice when decision makers are provided with the conditions necessary to calculate rationally (Cabantous et al. 2008), where those conditions may be constructed in organisations through the “collective mobilization of *theories* (economics and decision analysis), *actors* cognitively embedded in this theoretical framework, and *tools* designed to perform decisions in accordance with this theory” (Cabantous and Gond 2006, p 5, emphasis in original). For instance, Callon (1999, p 190-192) used the study of a strawberry market in the Sologne region of France to illustrate how the conditions for rational calculation by buyers and sellers were embedded within the physical market place and its procedures, the design of which was grounded in the principles of neoclassical economic theory. In this view, the rational decision maker is “formatted” (Callon 1998b) by a progressive translation of rational choice theory into standardised conventions for good decision making, and tools, methods, procedures, and techniques for decision analysis (Cabantous et al. 2008):

The rational decision making tools... contribute to the reconstruction of rational decision making within organizations by putting managers in a context similar to the one described theoretically by the rational choice approach of decision making. They embed theories and form institutional arrangements that frame decision makers' behaviours... They help turn the principles of “normative” rational decision making into a social reality directly accessible and potentially useful for decision makers... (*excerpt from Cabantous et al. 2008, p 412*)

It has been argued that world-level ERM frameworks such as COSO, AS/NZS 4360, and ISO 31000, can be seen to constitute standardised designs for “formatting” the rational management of uncertainty in organisations (Power 2004, 2007). Since it is the Chief Risk Officer who is tasked with implementing those designs in organisations, the CRO can be seen as normatively positioned to play the role of a Decision Engineer (March 1978) who formats rational decision making by embodying the theoretical conditions for rational choice within the design of ERM processes and infrastructure; thereby carrying out an important translation between theory and practice (Cabantous et al. 2008). The proposals put forward by Watercare’s CRM for developing the company’s risk management capabilities, and his focus on the quality of the company’s “risk data” in particular, reflected an interpretation of his role in these terms (see Chapter 3).

Chapters 4 and 5 have already called this approach into question based on an analysis of Watercare’s Corporate Risk Manager in action. He initially interpreted his role in terms

consistent with this strategy, but found it difficult and problematic to implement, and ultimately unproductive. It has also been argued in this chapter that it is problematic to talk about the task of decision support in terms of approximating the conditions of the so-called perfectly rational choice. Bearing in mind that “to format” means literally “to frame up” or “to arrange”, the strategy outlined above implies that the Decision Engineer would literally have to know all of the decision elements and how they should be framed, and be able to either provide this information to the actors in question, or be able to place those actors into a context where they could rapidly create the appropriate frame themselves. But such a strategy would only be possible if the Decision Engineer was an expert with regard to the particular problem or decision in question (indeed, one cannot provide expert guidance unless one is an expert). He or she would require not only subject expertise, but also expertise on the particular practice context as well, since it is only through such prior knowledge and experience that the Decision Engineer could know which elements should be taken into account and how they should be framed.

Unfortunately, the requirement of expertise makes it unreasonable to expect CROs to be able to employ this strategy in their roles, at least for the majority of the time. Chief Risk Officers may be experts in various fields, although most likely these would be insurance or finance-related, and may or may not have context-specific expertise depending on the length of their tenure in a particular firm, and experience in a particular industry. But, more importantly, CRO’s are tasked with organising Risk Management on an enterprise-wide basis. As such, CRO’s do not, and could not reasonably be expected to have subject- or context-specific expertise relevant to *all* of the contexts in which they may support decision making across the enterprise. In other words, in most contexts, CROs will probably know less than the actors they are trying to help about the particular problems or decisions that they are trying to frame. This was identified as a fundamental dilemma of the role in Chapter 5.

With reference to the earlier discussion, this also points up that the strategy of “formatting” or “pre-framing” can only realistically be applied in contexts of relatively low uncertainty and novelty. Decision making in such contexts may be made routine through standardised procedures, and constrained or regulated through the design and use of standardised decision rules, tools, and infrastructure. The strategy is not likely to be as successful in contexts of greater uncertainty and novelty where the decision elements are unknown in some significant sense.

## Strategy 2: Providing resources to the decision maker

If our would-be Decision Engineer knows less than the actors he is trying to help, if he is a generalist who does not already know what the decision frame should look like, then how is he to support decision making by those actors? What does “decision-support” mean in contexts of high novelty and uncertainty where no one is an expert and where there are no ready-made frames, but where, rather, decision frames are “in-the-making”?

Drawing on the earlier discussion, one possible strategy is that the Decision Engineer might provide “support” by manipulating the various contextual factors which constrain rational calculation (see the factors listed in Box 6.1 on page 190). He or she might do this, for instance, by providing additional information or resources, providing access to experts, bringing together relevant stakeholders, or by providing the time and space for communication and reflection. This is a more plausible strategy by which CROs might support organisational decision making since CROs are likely to have certain resources at their disposal and a degree of authority to manipulate the decision making circumstances of local actors. But if the Decision Engineer’s role is limited merely to providing resources to local actors then what is the legitimacy for such a role in organisations? Why not simply allocate such resources through the existing management structure of the firm? If the role is to have any legitimacy whatsoever then the Decision Engineer must be capable of supporting local actors in a way which does not merely duplicate the existing management responsibilities within the firm.

## Strategy 3: Helping decision makers to settle their frames

The third decision support strategy lies, conceptually, in the middle ground between the two approaches described above. That is, rather than defining (formatting) decision frames in the guise of an expert, or merely allocating resources to decision makers, the Decision Engineer may help actors to settle their decision frames by facilitating the identification and framing of various sources of uncertainty (or what were earlier referred to as the relations which overflow decision frames). The generic sources of uncertainty which may exist in any decision frame were discussed earlier and summarised in Table 6.7 (reproduced below). Each of those sources of uncertainty constitutes possibilities for framing the world, or a particular part of it, in different ways. The purpose of this third decision support strategy is, therefore, to help agents to identify and account for those possibilities by examining the sources of uncertainty with respect to their decision making endeavours.

*Copy of Table 6.7. Types and sources of epistemic uncertainty in decision frames*

<b>Framing distinctions</b>	<b>Type and source of uncertainty</b>
Recognition	<i>Ambivalence about the objective nature of things in the world, and about what will happen in the future.</i> This type of uncertainty arises from the extent and reliability of information and knowledge about the problem domain, and about the calculations that need to be made in that domain (including experiential knowledge).
Significance	<i>Equivocality about the value of things in the world.</i> This type of uncertainty arises due to bias, caprice, or conflict in the formation of value judgements
Relevance	<i>Mutability of relations between things.</i> This type of uncertainty arises due to errors or inconsistencies in the way things are arranged in relation to each other for the purposes of calculation.
All	<i>Fallibility of cognition.</i> This type of uncertainty arises from the reliability of agents' information processing capabilities (i.e. is the agent a reliable calculator or is the agent prone to making various sorts of errors such as bias, assumptions, mistakes?)

Since the stated purpose of Strategy 3 is essentially the same as the purpose of a Risk Assessment, it is necessary to clarify how it is in fact different. I have already described the key features of the Risk Assessment process as it is conventionally conceived in normative Risk Management models (see Chapter 3 and Appendix V) so there is no need to repeat them here. The following points of similarity with Strategy 3 may be noted:

- The purpose of a risk assessment is to identify, assess, and evaluate sources of uncertainty with respect to system performance. This requires: (i) a definition of the components and relationships for the system in question, and (ii) a methodology for placing those components and relationships under question so as to reveal the existence and significance of uncertainties in that system definition.
- The purpose of Strategy 3 is to identify, assess, and evaluate sources of uncertainty with respect to a decision or problem frame. This requires: (i) a definition of the entities and relationships which constitute the frame in question, and (ii) a methodology for placing those components and relationships under question so as to reveal the existence and significance of possibilities for alternative framings.

Risk Assessment seeks to question an agent's state of belief with respect to their definition of a particular system and expectations about its future performance, while Strategy 3 seeks to question an agent's state of belief with respect to their framing of a particular situation. To the extent that an agent's problem/decision frame can be interpreted as a physical system definition, and vice versa, then Strategy 3 and the generic Risk

Assessment process can be seen as conceptually equivalent. However, as it is imagined here, Strategy 3 constitutes a different vision of Risk Assessment than that depicted in the various normative models of Risk Management in the international literature.

The strategy is for CROs to intervene in organisations to interrupt the automatic flow of distinctions which make up decisions; in essence to bring those distinctions into question and force them to become matters for conscious reflection by organisational actors. Conceptually there are three opportunities to do this for any decision:

- i) Placing past decisions under audit provides an opportunity not just for oversight but also for learning from past performances;
- ii) Bringing current decisions into question provides an opportunity to examine and reflect on frames as they are constructed;
- iii) Bringing future decisions into question provides an opportunity to anticipate how future performances might play out.

Since it is impractical to examine decisions as they are being made in contexts of skillful action (i.e. it would not be reasonable to ask the race car driver to reflect on how he is making decisions while he is racing), only the first and latter of these approaches would be suitable as decision support strategies for such contexts. All three approaches would be suitable for situations of significant uncertainty and novelty where agents might experience difficulty framing their decision, but would also be suitable in situations of low novelty and uncertainty, often characterised by highly proceduralised or habitual decision making (Ryle 1949; Schermerhorn 1996). Placing such decision making processes under question provides an opportunity to identify where improvements might be made (e.g. breaking old habits, redesigning routines and procedures, etc.).

In each case, three key constraints of the CRO role orient the methodology for Strategy 3. First, acting as Decision Engineers, CROs cannot force other agents to settle on any particular decision frame. Whether and when an agent comes to hold beliefs or intentions, and what they may be, are all up to the agent. Framing is, in this sense, very much a personal endeavour. The Decision Engineer cannot force an agent to believe or act one way or another, but can only aid the agent in formulating those beliefs or intentions (this can be seen as a further limitation of Strategy 1 outlined earlier, i.e. even an expert must convince others that he is right). Second, the approach should not require the CRO to be an expert with respect to the particular problem or decision in question (nor present him as one). Rather, CROs should have the means and the legitimacy to engage organisational actors in certain times and

places and in ways which provide those actors with opportunities to productively question and reflect on their decision frames. And, third, CROs are unlikely to be in a position to actively provide decision makers with all of the time, information, and calculative resources that they need or might desire to completely resolve uncertainties. This indicates that the focus of the CRO's interventions should be on the *initial* identification and prioritisation of uncertainties, rather than their detailed analysis and calculation.

Taking these factors into account, an appropriate approach would resemble the knowledge facilitation methodology discussed in Chapter 4 (Roth et al. 2004). Very briefly, knowledge facilitation is conceptualised as a methodology for the creation and management of new knowledge in organisations (Alvesson and Kärreman 2001; Roth et al. 2004; Tsoukas and Vladimirou 2001). The function of the knowledge facilitator role is to create spaces where such learning processes can take place: "...the knowledge facilitator role is explicitly aimed at orchestrating the distribution of knowledge within the firm... to create arenas and opportunities... for sharing knowledge" (Roth et al. 2004, p 200; see also Chapter 4).

In this regard Strategy 3 is grounded in a paradigmatically different view of the Chief Risk Officer's role than that embodied by conventional notions of Risk Assessment. As it is conventionally conceived, Risk Assessment reflects a view of Risk as something real – i.e. the quantifiable uncertainty (probability) of something undesirable happening – which can be represented as Likelihood x Consequence. Risk Assessment is accordingly conceived as a more or less formalised methodology for assessing Risk, characterised by the normative ideals of explicitation, specificity, and quantification. In this view the goal of decision support interventions is to ensure that decision making is informed by good "risk data", which can be achieved through the formal assessment and calculation of Risk. This conventional view is summarised in the left-hand column of Table 6.9.

This thesis has substantially challenged the efficacy of this paradigmatic view of the CRO role. In contrast, the decision support strategy outlined here – i.e. Strategy 3 – reflects a broad view of risk as merely a label for sources of uncertainty in knowledge (where the focus of the term "uncertainty" is primarily on the epistemic aspects of uncertainty). It acknowledges that, depending on context, actors may have more or less specific perceptions of those sources of uncertainty. The assumed goal of decision support is to help agents make decisions which are reasonable; that is, consistent, and accommodating of this uncertainty. But since decisions are necessarily personal this goal cannot be pursued in a prescriptive manner. The decision support strategy is therefore conceptualised as a convened function in organisations, the purpose of which is to facilitate actors' reflection about the distinctions which constitute decision frames in order to reveal potential uncertainties. This alternative



view of the CRO role is summarised in the right-hand column of Table 6.9.

*Table 6.9. Two paradigmatic views of the Chief Risk Officer's role in decision support*

<b>Key objects of attention</b>	<b>Conventional view</b>	<b>Reconceptualised view</b>
Risk (Uncertainty)	Risk is something real: quantifiable uncertainty (i.e. probability) of something undesirable happening; can be represented as Likelihood x Consequence.	Uncertainty refers to sources of uncertainty in knowledge in a primarily epistemic sense. Actors may have more or less specific perceptions of these sources of uncertainty.
Risk Assessment (Strategy 3)	Risk assessment is a formal methodology or procedure to assess Risk, characterised by explicitation, specificity, and quantification.	Strategy 3 is a convened function within organisational planning cycles which seeks to facilitate awareness, understanding, and prioritisation of uncertainties of attention based on an agreed vision of the world
Decision support role of the CRO	Good decisions should be based on good data. Therefore the goal of decision support is to produce good "risk data", which can be achieved by formal assessment and calculation of Risk.	Good decisions are reasonable, but decisions are also personal. Therefore facilitate reflection about distinctions which constitute decision frames in order to reveal potential uncertainties.

### **Strategy 3 as an integral planning function in organisations**

From the latter perspective, above, Strategy 3 is envisioned as a methodology by which CROs can support organisational decision making by providing actors with opportunities to productively question and reflect on how they frame their worlds. It necessarily requires the CRO to intentionally and explicitly engage actors in certain times and places, and for this reason constitutes a disruption of the day-to-day activities in the organisation. The challenge for CROs is to embed the process as a normal part of management activity at any level in the organisation.

The framework in Table 6.6 (p 203) provides a basis for conceptualising how this might be achieved. While any agent may fulfil any or perhaps all of the problem solver roles in Table 6.6 depending on their particular relationship with their problem solving context, there is no reason why the framework in Table 6.6 must be interpreted only in terms of individual human agents. The framework may be interpreted from an Actor-Network perspective, wherein problem solving entities are conceptualised as heterogeneous networks (systems) of human and non-human actants which performs various kinds of problem solving functions (Callon 1998b, 1998a; Callon and Latour 1981; Latour 1987, 2005; Law 1992, 1999; Law and Singleton 2000). From this perspective, the problem categories, roles, and functions are

fulfilled by such problem solving entities, of which individual human agents are only one possible type. This is similar to ideas about distributed cognition, in which the intelligent processes of human activity are modelled as transcending the boundaries of the individual actor (Hutchins 1995). Cognition, learning, and the production of novel knowledge are re-interpreted as processes occurring through systems of heterogeneous (human, technical, social) components. It also resonates with theories of the evolution of Large Technical Systems, in which the problem solving system encompasses not just the technical artefacts of the system, but also the people and organisations which develop, operate, and maintain the physical infrastructure, the scientific (knowledge) artefacts that they employ, and the institutional arrangements (cultural, political, legal, and economic) in which they are embedded (Hughes 1987; Joerges 1988; Monstadt and Naumann 2005; van der Vleuten 2004). Thus, workgroups and teams, business units, and entire organisations can be seen as collective problem solving entities, and it is not just the human agents within those entities who perform the problem-solving work. Rather, the activities of framing and calculating are performed in a distributed fashion by and through the network as a whole.

From this perspective, organisations are constituted by groups and systems which are concerned with the management of certain types of problems. That is, organisations are created and sustained to fulfil the purposes of their owners (and primary stakeholders), and the fulfilment of those purposes requires their translation through Strategic, Planning, and Operational levels, as depicted in the framework in Table 6.6. But the labels of Strategic, Planning, and Operational are now defined relative to the organisation as a whole, and the translation of purposes is understood as being performed not just by human actors, but rather by the heterogeneous networks of human and non-human actants which constitute the organisation. This is illustrated in Table 6.10 which maps the generic problem types from Table 6.6 against the generic business processes of Operations, Planning, and Projects, based on the nature of the objectives those processes are intended to fulfil.

Of the three opportunities for intervention mentioned earlier, it is envisaged that the primary focus for CROs will be to bring current and near future decisions into question. This positions Strategy 3 as a process intended to support the organisation's planning (resource allocation) activities in each of the above contexts (to the extent that the methodology is employed to audit decisions after the fact it serves as learning process which informs future planning activities). The objective for CROs would be to identify when and where Strategy 3 interventions should take place within an organisation's core business processes, and what form those interventions should take.

Two critical understandings, already mentioned, must frame this work. First, CROs must

understand that their role is to facilitate knowledge-sharing and reflection on problem and decision frames, not to dictate what those frames should be. Second, the objective of the process is the *initial* identification and prioritisation of uncertainties, not their detailed analysis and calculation. It is the responsibility of organisational actors to subsequently perform the necessary investigations, analyses, and calculations as part of taking action to manage the identified priorities. The corollary of these points is that the appropriate mechanism for considering cross-functional and enterprise-wide issues in any particular instance will be to involve the right people, either by “getting them in the room” or facilitating their attention to the outputs of the process. This points up the importance of the CRO’s attention to who should be involved (see Table 6.6, p 203).

The vision is that Strategy 3 should become an integral part of planning activities at any level in the organisation (see Table 6.11). When this is achieved, organisational actors become aware of when and where the process will be performed, they prepare for it as a normal planning activity, they commit to take action on the identified priorities in their contexts of action, and formal records of identified priorities do not reside with the CRO as stand-alone “risk registers” but rather constitute the working planning documents of the business, e.g. asset management plans, project delivery plans, business plans, strategic plans. Together with other important texts (e.g. Constitutions, Statements of Intent, Annual Reports, policy statements, etc.), those documents should constitute a live and evolving vision of the organisation and its context, its objectives and priorities, and its performance<sup>17</sup>.

**Table 6.10.** *Business processes address different types of problems*

<b>Process</b>	<b>Objectives</b>	<b>Problem type (Table 6.6)</b>
Operations	Ensure the quality, timeliness, and cost-effectiveness of product and service delivery.	Operational problems
Planning	Determine what needs to be done to maintain and improve existing capabilities.	Planning problems
	Determine where and when new capabilities are required.	Strategic problems
Projects	Determine what needs to be done to deliver the project.	Planning problems
	Ensure the quality, timeliness, and cost-effectiveness of project delivery.	Operational problems

<sup>17</sup> This points up the duplicity of vision involved in developing standalone risk frameworks and registers.

**Table 6.11.** Conceptualising Strategy 3 as planning support in different contexts within the firm

<b>Context</b>	<b>Strategy 3 is a process/methodology/procedure for:</b>	<b>Problem type (Table 6.5)</b>
Planning (e.g. annual business planning, strategic planning, asset management planning)	Engaging and facilitating organisational stakeholders to: <ul style="list-style-type: none"> <li>– Examine what the business (or part of it) already does, in the context in which it does it;</li> <li>– Identify possible threats to continued reliable performance, possible opportunities to improve performance, and potential actions to address those possibilities;</li> <li>– Prioritise those actions for attention and resource allocation.</li> </ul>	Strategic Planning
Operations (including project delivery)	Engaging and facilitating organisational stakeholders to: <ul style="list-style-type: none"> <li>– Examine what they are about to do in the context in which they are about to do it;</li> <li>– Identify possibilities for error and improvement, and actions to take to address those possibilities;</li> <li>– Orient their attention (awareness) to the critical possibilities.</li> </ul>	Operational
Project Planning	Engaging and facilitating organisational stakeholders to: <ul style="list-style-type: none"> <li>– Examine proposed changes to business operations in the context of the enterprise as a whole;</li> <li>– Identify possible threats and opportunities with respect to both the project and the project's effect on the enterprise, and potential actions to address those possibilities;</li> <li>– Prioritise those actions for attention and resource allocation.</li> </ul>	Planning Operational
<p><u>Notes:</u></p> <p>The focus of the process is the <i>initial</i> identification and prioritisation of uncertainties, not their detailed analysis and calculation.</p> <p>Strategy 3 may take on a variety of forms depending on the particular problem context (right-hand column). For instance, a well-designed checklist coupled with an ingrained culture of awareness may be sufficient to implement Strategy 3 in certain operational contexts, whereas a series of facilitated workshops supported by behind-the-scenes co-ordination and collation may be more appropriate to support strategic planning processes.</p>		

### Generic vision of the methodology for Strategy 3

In order to put Strategy 3 into practice in any context, CROs require a methodology for intervention in three phases: (i) Establishment, (ii) Opening up, and (iii) Closing down (Sheffield 2005). The general requirements of these phases of the methodology are outlined in the following sections.

#### Methodology Phase 1: Establishment

Conceptually, the Establishment phase of the methodology is intended to establish arenas and opportunities for agents to make explicit, call into question, and reflect on their individual and collective framings of the world (or the particular situation in question). These arenas and opportunities need to be created in three dimensions: spatially, temporally, and organisationally (culturally).

- Spatial dimension: Although it might seem obvious, actors need a physical space within which they can engage in the reflective performance. This may be nothing more than a standard meeting room with tables, chairs, and a whiteboard, or it may require more elaborate design (e.g. a designated retreat spatially differentiated from the usual place of work, or even a specially tailored environment designed to provoke questions and reflections).
- Temporal dimension: Again, it might seem obvious, but busy actors need time to reflect in productive ways. Opportunities to do so might be provided within dedicated workshops or retreats, or they might also be created and embedded as a routine practice within the actors' normal work environs.
- Organisational (cultural) dimension: In Table 6.6 the problem solving roles that agents fulfil were defined in terms of the translation of interests and the transmission of information. In an organisational sense, this calls attention to the policies, procedures, and politics which define and constrain the contexts within which people work. If agents are to reflect in productive ways (thinking outside the box, imagining new possibilities), then they also need to be given the freedom and flexibility, at least temporally and within reason, to escape those constraints. This might be achieved, for instance, through an absence of power relations in the room, and facilitated engagements which give everyone fair opportunity to speak. With reference to the points made below about the Opening Up phase of the methodology, the establishment of trust between actors is likely to be an important part of the

Establishment phase.

The outputs of Phase 1 of the methodology are: *The time, space, and freedom to work together productively are created, and trust is established between participants.*

### **Methodology Phase 2: Opening Up**

Conceptually, the Opening Up phase of the methodology might be aligned with the Context Definition and Risk Identification phases of the generic Risk Assessment process. The objective of the Opening Up phase is to examine the four sources of uncertainty identified in Table 6.7. In order to do this a certain amount of “making explicit” will be required – the world or system or process which is to be examined and reflected upon must be made explicit for all to see. This means that what is known about entities and their relations, and agents’ judgements and expectations must be made explicit; to the extent that they remain implicit, it will be that much harder to reveal the locations of uncertainties.

Since uncertainty, equivocality, and subjectivity are all a function of the strength of the actor-networks which support and sustain knowledge claims, judgements, and expectations (see Chapter 2), in order to call frames into question, one must open up those networks. There are only two ways in which this can be achieved:

- Deconstruct the network: that is, deconstruct knowledge claims, judgements of relevance and significance, and experience-based expectations. In Latour’s (1988) terms this is to subject those networks to “trials of strength”, to find out what supports them, and in doing so to reveal the weaknesses and unaccounted-for relations in those networks. However, the deconstructive task is likely to be both difficult and costly, since it requires expertise and infrastructure of a scale and magnitude comparable to that which constructed the network in the first place (Law and Singleton 2000).
- Bring in credible witnesses: In the absence of the expertise and infrastructure necessary to deconstruct the knowledge claims, judgements, and expectations expressed by other agents, one must trust those agents as reliable bearers of information. This emphasises the importance of building trust (credibility) between actors, particularly in situations involving an extended stakeholder community (Yankelovich 1999). It also highlights that in these situations the only way of opening up and challenging frames is essentially by proxy, i.e. by bringing in other credible actors who do have the necessary expertise or political support to mount a “trial of

strength”.

It is foreseeable that this process “opening up” may well be one of revealing and reconciling competing views to establish an agreed vision of the world. That is, different actors will bring different knowledge, experiences, and values into the room, all of which will influence how they see the world. The greater the disparity of stakeholders present, the broader the range of world-views will be. This points up a key feature of the methodology. Where multiple stakeholders are involved, and where those stakeholders represent different communities, the knowledge, values, and expectations brought to the table will be very different. In such situations it is unreasonable to expect those actors to be able to develop confidence in the claims of others based solely on their ability to deconstruct those claims. Rather, that confidence will necessarily be a product of practical judgements about the credibility of those other actors. That is, each actor will have to judge a given claim against the following question: Do I trust the person or group presenting the claim? This highlights the importance of building trust and facilitating effective communication between stakeholders, and of getting the right people in the room.

The outputs of Phase 2 of the methodology are: *An agreed vision of the world is established between stakeholders, and the key sources of uncertainty underlying that vision are identified.*

### **Methodology Phase 3: Closing Down**

Conceptually, the Closing Down phase of the methodology could be aligned with the Risk Analysis and Evaluation phases of the generic Risk Assessment process. However, the intention is that this work should be performed at a relatively coarse level. I have already mentioned that CROs are unlikely to be in a position to actively provide decision makers with all of the time, resources, and calculative resources that they need or might desire to completely resolve uncertainties. Chief Risk Officers are also unlikely to be in a position to facilitate or provide guidance for such detailed calculative work. Therefore, the objective of the Closing Down part of the methodology should be to facilitate the preliminary evaluation of the relevance and significance of identified uncertainties. This is necessary in order to establish priorities for further work by actors going forward.

Since the process is unlikely to have productive outcomes without the personal commitment of the involved actors to the resultant evaluations and priorities, a critical aspect of the methodology will foreseeably be to facilitate that personal commitment (Sheffield 2005). Successful integration of the methodology as a core planning function should be conducive in this regard.

The outputs of Phase 3 of the methodology are: *Priorities for action going forward are agreed, and the commitment of individual actors to address those priorities is established.*

### **Variations to the methodology**

The framework in Table 6.6 identifies the factors which are likely to differentiate the intervention methodology for Strategy 3 across problem-solving/decision-making contexts. That is, the methodology for supporting decision makers in the three types of context (Operational, Planning, Strategic) may need to be varied to account for differences in:

- The uncertainty and novelty perceived by actors, and difficulty experienced with framing and calculating;
- The number and range of stakeholders and interests involved in the process;
- The relative significance of values-based (axiological) versus knowledge/information-based (epistemological) uncertainties;
- The likelihood of contest and conflict between perspectives and interests in both the axiological and epistemological dimensions;
- What is uncertain each context (e.g. objectives and decision criteria versus methods and options versus the accuracy and precision of measurements and calculations).

However, it is beyond the scope of this project to attempt to map out in detail what that methodology might look like in specific contexts, and what specific methods and tools CROs should use in different situations. In this regard, the work in this chapter may serve as a stepping off point for future research into and development of specific methodologies, methods, and tools. Such work may well draw on and adapt a range of extant methods and tools which have been developed and applied for similar purposes in other fields (including, for instance, risk assessment, scenario planning, facilitation, strategic planning).



# Chapter 7

## Conclusions

Chapter 1 identified that two decades after the emergence of Enterprise Risk Management and the Chief Risk Officer role on the international stage, the knowledge base for performing the role of CRO, as distinct from managing risk, remains significantly underdeveloped. Specifically, there is a dearth of explanatory theory and a lack of insight with respect to translating world-level concepts of ERM into operational systems and processes in organisations. This thesis has contributed to filling this gap.

The thesis has explored ways in which Chief Risk Officers may understand and approach the decision support function of their roles. Empirical grounding was provided by a detailed, longitudinal case study of the implementation of Enterprise Risk Management by the Corporate Risk Manager at Watercare Services Ltd. The case study was significant because of the CRM's principal focus on how ERM could support "better" decision making in the organisation. Theoretical grounding was provided by drawing on literatures from a range of disciplines to interpret and analyse the particulars of the CRM's performance, and to theorise the decision support function of the CRO role. The following question served as the focal lens for the inquiry: How should CRO's understand their role with respect to "improving" organisational decision making, and what strategies might they employ toward this goal? This final chapter reflects on the substantive conclusions in respect of this question, the contributions that the thesis makes to the CRO knowledge base, and the directions in which future research might extend and develop those contributions.

Chapter 1 also identified that the nature and contexts of engineering work are changing in the 21st century. Professional engineering work is becoming increasingly distributed, multi-cultural, cross-disciplinary, and complex in nature. Engineers have to be able to co-

ordinate and work within cross-disciplinary teams, and formulate solutions to complex problems with antecedents and consequences across ecological, economic, political, legal, and cultural domains. While technical know-how remains the core of the engineering knowledge base, it is the broader professional skills which are increasingly valued in the workplace, leaving pure technical specialisation as a mass commodity. Since the late 1990s a growing discourse has argued that engineering education must evolve to better prepare young engineers for the realities of professional practice in the 21<sup>st</sup> century (see Chapter 1). As a result of its particular positioning, the thesis has generated insights relevant to the practice and study of engineering work, particularly where engineers step into cross-disciplinary roles which challenge the limits of the conventional engineering education. This chapter therefore also reflects on how the thesis contributes to the future development of the engineering profession.

## Conclusions

### **About the design of risk frameworks and registers to support organisational decision making**

*Risk frameworks should reflect “what matters”, but the answer to the question of “what matters most” will depend on who is answering it.*

Watercare’s Corporate Risk Manager discovered that there was an inherent subjectivity to all the decisions involved in the design of the risk framework: decisions about which objectives to represent, which performance parameters or criteria to use, about the relative priorities of those objectives and parameters, and about how to organise them into a coherent framework. Contrary to the CRM’s initial assumption, these decisions could not be reached via a detached deconstruction of the corporate objectives and performance standards contained in various texts within the organisation. Rather, in order to specify a framework which did accurately represent priorities throughout the organisation the CRM had to engage with various stakeholder groups to uncover and make explicit perceptions of “what mattered” in specific functional contexts (i.e. general managers in the strategic context, project managers in the project delivery context, and operations managers in the operations context). The CRM referred to this process as being “more art than science”.

*While detailed risk frameworks and registers appear to fulfil aspirations to objectivity, they may, on the whole, be unproductive.*

Epistemologically, risk is not just a claim about certain things in the world but is also a state

of knowledge about those things. That is, we can talk about risk when we know enough to reasonably calculate and specify the real probabilities of outcomes. Where we cannot make that estimate then we are merely uncertain. Detailed risk frameworks and registers appear to fulfil aspirations to the objectivity of “risk data” because they require users to both possess and express a deeper and more refined knowledge of risks. However, as the CRM’s experience showed, they are also difficult, time consuming, and costly to produce, and for this reason are foreseeably inflexible in the face of changing objectives and performance standards. Furthermore, since the necessary distinctions only become possible toward the end of what might be a fairly lengthy and involved process of inquiry this raises problems for the productive use of detailed risk frameworks and registers in organisations. Such infrastructure is not productive in the early stages of inquiry which are marked by uncertainty, ambiguity, and vagueness about both means and ends (this was illustrated by the difficulties experienced with the CRM’s new project risk register). And not all inquiries produce such detailed “risk data”, in which case such infrastructure either drives more costly risk assessments than might otherwise be required, or risks the misrepresentation of “risk data” as being more objective than it really is. This suggests that in addition to being appropriate to the functional contexts of the organisation, ERM infrastructure should also be appropriate to the knowledge production contexts in which it will be employed, if it is to be used effectively. But this implies a significant practical burden on CROs, not only to understand the needs of actors in different functional contexts across the organisation, but also to understand the needs of actors engaged in different phases of the knowledge production process and which tools will be appropriate to support those needs.

*The normative goal of objectivity must always encounter the inevitable practical constraints to its achievement.*

With respect to the design of risk frameworks and registers, the overarching value problematic can be framed as follows. Taken to its ultimate extreme, the goal of an objective risk framework would amount to representing all of the organisation’s objectives, performance standards and parameters in absolute detail, and in a format that made explicit all of the relative priorities between those objectives and standards. In other words, the risk framework would have a 1-to-1 correspondence with the company’s metrological performance system. But, given the complexity of that system for most organisations, such a framework would be utterly impractical both from a design perspective and in use. The other end of the spectrum is equally undesirable, where the correspondence between the risk framework and the company’s objectives and performance standards would be so loose as to

promote rather than constrain subjectivity in use. Conceptually, a suitable risk framework must lie somewhere between these two extremes.

*Practically, the design of risk frameworks and registers involves reconciling a number of objectives which are not necessarily commensurate.*

The risk framework should: facilitate the efficient and accurate communication of risk perceptions between different audiences in the organisation; while providing a consistent and coherent basis for the capture of risk information; but without imposing an unnecessarily costly bureaucratic or calculative burden on the organisation; or creating a situation where the objectivity (accuracy, fidelity, reliability, etc.) of that information is misleadingly represented. The CRM's experience suggests that neither "one-size-fits-all" nor "tailor-to-context" approaches can reasonably satisfy all of these objectives at once. Indeed, in light of the discussions in Chapters 5 and 6 it seems likely that these objectives cannot be satisfied through the design of ERM infrastructure alone. Rather, CROs should consider how those objectives might be fulfilled through a range of structural and non-structural (i.e. facilitative, process-based) interventions.

### **About how CROs should understand and approach the decision support function of their role**

*There are different legitimate ways of conceptualising Enterprise Risk Management and thus of interpreting the role of the Chief Risk Officer. Those different ways of "seeing" imply different possibilities for acting, which may be more or less effective for achieving desired outcomes.*

First, the role may be interpreted in different ways depending on how the functions of ERM are emphasised; i.e. whether emphasis is given to supporting and enhancing organisational decision making (the decision support function), or to supporting internal control and providing assurance of organisational processes through various procedural and administrative technologies (the compliance function). Second, the decision support function of the role may be approached in different ways depending on how one perceives Risk and the purpose of Risk Assessments. For instance, as it is conventionally conceived, Risk Assessment reflects a view of Risk as something real – i.e. the quantifiable uncertainty (probability) of something undesirable happening – which can be represented as Likelihood x Consequence. Risk Assessment is accordingly conceived as a more or less formalised methodology for assessing Risk, characterised by the normative ideals of explicitation, specificity, and quantification. In this view the goal of decision support interventions is to ensure that decision making is informed by good "risk data", which can be achieved through

the formal assessment and calculation of Risk. Watercare's Corporate Risk Manager initially conceived his decision support interventions in these terms. Alternatively, the decision support approach may be conceived more broadly if risk is understood as merely a label for sources of uncertainty in knowledge (where the focus of the term "uncertainty" is primarily on the epistemic aspects of uncertainty). It acknowledges that, depending on context, actors may have more or less specific perceptions of those sources of uncertainty. The assumed goal of decision support is to help agents make decisions which are reasonable; that is, consistent, and accommodating of this uncertainty. But since decisions are necessarily personal this goal cannot be pursued in a prescriptive manner. The decision support strategy is therefore conceptualised as a convened function in organisations, the purpose of which is to facilitate actors' reflection about the distinctions which constitute decision frames in order to reveal potential uncertainties.

*The latter of the above approaches to decision support, which I called "framing facilitator", is more likely to lead to productive outcomes than alternative approaches.*

It is possible that CROs might choose to support organisational decision making by manipulating or influencing the various contextual factors which constrain rational calculation. But to the extent that this approach is at root always an issue of resource allocation, it merely supplements existing management responsibilities and is therefore only weakly legitimate. A more direct form of support would be to provide expert guidance to decision makers. However, this approach encounters the limits of expertise in at least three ways. First, to the extent that it is associated with aspirations to analytical objectivity and calculative precision, it will be constrained by the capabilities of organisational actors to specify and calculate Risk in accordance with its theoretical specification. Second, CROs are generalists and will therefore be constrained in their ability to provide expert guidance in any and every context in which risks must be analysed and managed (the particular backgrounds of individual CROs notwithstanding). Third, in the absence of such expertise, CROs are likely to encounter difficulties convincing other expert actors that they need to develop and use more advanced methods and tools for knowing risk. These latter two constraints constitute what I referred to as the CRO's dilemma (see Chapter 5). While CROs might successfully employ either or both of these two strategies in various contexts, the thesis has argued that the "framing facilitator" approach is more appropriate as a general decision support strategy for CROs because it is cognisant of the key constraints on the CRO position and is grounded in key understandings of decision making behaviour.

## **About developing the knowledge base for professional engineering**

*The study presented through this thesis is an example of how ethnographic studies can contribute to the development of the professional knowledge base.*

This thesis has presented a study of a professional role, and contributed toward the explication of a theory of action for that role. The subject of the study may not have been an engineering role, even if it was performed by an engineer in an engineering firm, but it may nevertheless be seen as an example of how the close, ethnographic study of roles can contribute to the development of the professional knowledge base. Watson (1997, 2011), for instance, has argued for the value of ethnographic studies in management as providing critical insight into “how things work” in organisations and thus contributing to the development of the knowledge base for professional managers. It may therefore be suggested that the project to broaden the education of professional engineers, whether in undergraduate, post-graduate, or Continuing Professional Development contexts, could be productively informed by similar ethnographical studies of “how engineers work”, as distinct from what they work on or with. The shadowing of engineers in their day-to-day work would serve to illuminate the broader professional, managerial, ethical, and leadership work in which engineers engage in their roles, the kinds of the problems that they encounter in performing this work, how it relates to the more fundamental, technical aspects of their roles, and thus what they might need to know in order to perform their roles more effectively and efficiently. This might, in turn, inform the development and teaching of the formal knowledge base and methodologies for coping with the non-technical problems that professional engineers encounter; as this thesis has contributed to the professional knowledge base for the role of Chief Risk Officer.

*If engineers are to step into strategic roles as the leaders and facilitators of interdisciplinary problem solving teams, then they will require approaches and knowledges which transcend not only individual disciplines in engineering, but also the profession of engineering as a whole.*

There have been calls for a new type of engineer, envisioned as an integrative, transdisciplinary role, whose purpose is not to engage in design in the traditional sense, but rather to foster collaboration between engineers and non-engineers and to synthesise innovative design approaches to complex problems (see Chapter 1). This thesis presents a study of two engineers moving into just such strategic positions, one in a practice context, and the other in a research context.

In this regard the thesis presents a cautionary tale about the efficacy of the engineering

approach with respect to formulating productive solutions to the kinds of problems with which such “strategic” engineering roles are concerned. The central movement that occurs through the narrative of the thesis is a shift between two very different perspectives on how to “implement” ERM. The first perspective was a functional, structural approach where thinking about frameworks, hierarchies, structures, and calculation dominated. The other perspective is a process-based approach where thinking about knowledge production, emergence, interaction, and collaboration dominate. The movement between the two perspectives was an empirical one. Watercare’s CRM initially approached his role from the former perspective, belying his background as an engineer. But over the course of the research period, the CRM gradually evolved his approach to the latter perspective as a result of realisations which emerged through experience, dialogue, and the contrast between his conceptual and practical worlds. That is, at least initially, the CRM operated in a conceptual world where risk was expected to be something specific and well-defined that could, and indeed should be calculated with objectivity and precision. In practice, however, he found himself working mostly in a realm where things were subjective, ambiguous, and uncertain; quantitative, precise, clearly defined risk was rarely to be found, and, indeed, was only locatable with much effort. In this domain, the CRM’s carefully planned interventions to “design and build” risk frameworks and registers encountered significant resistances (Pickering 1993, 1994), and the CRM was forced to reconceptualise his approach in distinctly non-engineering terms.

Although the inquiry has been framed as a study of the implementation of ERM, it shows both the CRM and myself searching for the means to perform the kind of boundary-spanning work that was described above, and in this task having to reach beyond our own knowledges and experiences in engineering. Thus, it would seem logical to conclude that if and when engineers seek to step into the kinds of strategic, leadership roles that have been promoted, then they must be ready to encounter the limits of engineering knowledge and methodology. Indeed, such movements can be seen as stretching or even throwing into question the very category of “engineer”: that is, when an engineer no longer draws on knowledge of mathematics and the physical sciences to design things, but instead draws on knowledge of the social sciences help others to formulate and solve design problems, then is he or she still performing an engineering role?

*The thesis contributes to the development of transdisciplinary inquiry in engineering.*

Both the study of engineering roles (i.e. how engineers work) and the performance of “strategic” engineering roles, described above, require methodologies for transdisciplinary

inquiry. Chapter 2 justified a transdisciplinary methodology for the study of a professional role. Similar approaches could be employed to study what engineers do, how they do it, and what they need to know in order to do it, in the contexts in which they work. Chapter 6 justified a general strategy that CROs may employ to intervene in organisations to support decision making processes. The strategy may be applicable as knowledge-for-action for the role of Strategic Engineer.

## Contributions

In light of the aforementioned inadequacies in the relevant bodies of literature, the thesis has added greatly to knowledge and understanding, both of and for the CRO role. In order to properly locate the contributions of the thesis I differentiated what I called a general Theory of Enterprise Risk Management from what I called a Theory of Action for the Chief Risk Officer. The former conceptually demarcates a domain of knowledge and theory which describes and explains *firm-level effects* of ERM implementation, and correlates those effects with certain types of ERM interventions within the organisational milieu. The latter conceptually demarcates a domain of knowledge and theory which describes and explains (and hence informs) the *context-specific actions* of CROs. This latter body of knowledge is concerned with what CROs do, how CROs choose which interventions to pursue, and how they produce those interventions in specific contexts. The thesis has contributed toward such a Theory of Action for the role of Chief Risk Officer by drawing insights from the analysis of a detailed account of a CRO in action, and by theorising the decision support function of the role. The thesis offers the following substantive contributions:

- Insights and understandings about the design of risk frameworks and registers, and their efficacy as instruments for supporting organisational decision making;
- Corroboration of Mikes (2010) finding that the role of CRO may be conceptualised and performed in different ways. The thesis identifies and discusses two perspectives and approaches, which are characteristically similar to those identified and labelled by Mikes (2010) as “strategic controller” and “strategic advisor”. The thesis surfaces the constitutive assumptions behind the two perspectives.
- A theoretical framework describing a general typology of problem/decision situations, which may be used to interpret organisations as problem solving systems. The framework makes explicit the key general features of the range of organisational contexts in which a CRO might intervene in order to help agents make better



decisions.

- Theoretical justification of a generic strategy that CROs may employ to intervene in organisations to support decision making processes, which is grounded in fundamental understandings of decision making behaviour, and which is cognisant of the unique constraints on the CRO position.

The thesis is also relevant to engineers. First, where engineering firms adopt ERM the task of implementing that programme will fall to the engineers who run those enterprises (Lewin 2006). This thesis informs that task. Second, and more significantly, the thesis can be seen as a contribution to the current debate about the appropriate scope of education for engineers in the 21<sup>st</sup> Century. The thesis offers the following contributions in this regard:

- An example of how transdisciplinary research into professional roles can contribute to the development of the professional knowledge base. Similar approaches could be used to study how professional engineers work, as distinct from what they work on or with, and thus to orient and develop the formal curriculum for the broader education of engineers;
- Transdisciplinary methodology for both the study of engineering roles (Chapter 2), and for performing the kind of cross-disciplinary work that will characterise certain “strategic” engineering roles in the 21<sup>st</sup> Century (Chapter 6).

## Future Research

This research may be taken forward in a number of directions:

- The thesis largely ignores the internal control and compliance function of the role (i.e. that associated with the governance claim of ERM). While this latter function may well be the major component of CRO roles in most organisations, it was simply beyond the time and capacity of this study to theorise two core functions of the role, let alone their interrelation. The study of this function, and how it interrelates with the decision support function of the CRO role, would be a valuable contribution to the CRO knowledge base.
- The thesis also does not consider of the question of whose interests the CRO serves. Where the CRO is located in the corporate structure, who the CRO reports to, the nature of the institutional and regulatory environment of the organisation, the internal power structures of the organisation, and the CRO’s own personal and

professional values are all factors which will necessarily influence how individual CROs approach their roles. The effects of these and other contextual features of the role had to be excluded from the theoretical scope of this thesis due to the time and capacity limitations of the PhD process. Future research on these factors may further illuminate the complexities of the role.

- As regards what constitutes a “good” decision, the thesis considers only the criterion of rationality. There are, of course, other criteria against which the process and outcomes of decision making may be judged (e.g. objectivity, transparency, inclusiveness, fairness, etc.), but it was not possible to consider the implications of such a broad range of criteria within the limitations of this study. Future research may generate additional decision support strategies for CROs which accommodate these broader criteria.
- The development of the framework (typology of problem/decision situations) and the decision support strategy in Chapter 6 was based primarily on theoretical work. An important avenue for future research will be to evaluate their workability as guides for practice. In particular, the framework implies a methodology for interpreting the organisation as a problem solving system, but the thesis does not develop this methodology. Future research could contribute significantly to the CRO knowledge base by elaborating a general methodology by which CROs may interpret and understand key decision making processes in their organisations, and subsequently integrate the decision support strategy from Chapter 6 into those processes.
- The decision support strategy developed in Chapter 6 was elaborated only in very general terms. The CRO body of knowledge could benefit greatly from the extension and development of this strategy to include specific methods and tools that CROs might employ to support different kinds of decisions in different contexts. A range of methods and tools probably already exist in various disciplines, which might be borrowed for the CRO toolbox.
- The methodology described in Chapter 2 could be applied as a basis for shadowing engineers in action in order to record, describe, and theorise how engineers work. Such research could make valuable contributions to the broader professional knowledge base of engineers.

# Dialogues



# Dialogue No. 1 – The Vision

**Source:** Transcript 15 June 2007 (CRM:03)

**Participants:** Corporate Risk Manager (CRM)  
Richard Donnelly (RD) [*Researcher*]

---

CRM: In the last couple of weeks I've put a lot of thought into putting a stake in the ground as to where I want to take risk management here. I'm conscious of having been here now four months and I need to make some kind of progress. So I've been drafting a document on the vision for the business, writing up what I thought the risk management function should be able to deliver. So I've been thinking "what is the vision?" How do you define what the vision is? And I found myself concentrating on what mathematical things, what analysis I wanted the risk management function to be able to deliver. I guess I focussed on the concept that risk management is supposed to support decision making and being an engineer I went to that thing that decisions will be based on data. So how can we provide data in a format and a level of detail that will enhance the way that we currently make decisions? And when I read my risk management vision now I see that it's all about tools and analysis, and I realised that there was nothing in there about the culture of the organisation and about how risk management is practised. Actually that's interesting as to whether it's a practice or a support function, or whether it's there to deliver data to help you make decisions, or whether it's an action. Well I guess their obviously linked aren't they? You make decisions about how you're going to act. So anyway, when I went back through it I thought that it was weak on the culture side.

RD: What made you refocus on the culture and context side of it when you re-read the paper?

CRM: Ah...yes I know why. I was having a discussion with someone and it reminded me about the people in the business, that there's a need for training because some of them don't even feel comfortable with the current risk management function. So that reminded me that I needed to be doing training, and I realised that's not included in my vision. When I went back to add a section on training, I kind of made the association with culture, and realised that the document didn't deal very much with culture. Anyway, so I was focussed on what functionality I was trying to deliver, that I was trying to deliver decision support in areas where I think the decisions to the company are important. And that made me realise that I don't really understand where the important decisions in the company are. I still haven't felt like I've got a clear understanding of the business. So just yesterday I started going through the bulk water contracts and the Asset Management Plan, some of the primary

documents that the organisation either produces or that govern the organisation. I started reading them to try to determine what's important and what's not, what are the influential factors in the business.

RD: What sort of factors are you looking for?

CRM: I don't know. I'm still trying to get a handle on it. You know, after you've been in a business for a long time, or if you've been in any environment for a long time, you start to intuitively understand what's important and what's not. But I haven't got that level of understanding yet. The idea being that if I know what the top five factors are that really have an impact on the business then I could focus my development efforts on supporting those five factors. But to be honest it's been a little bit difficult trying to find out what they are. I could go and talk to, in fact I probably should go and talk to some other people, like the GMs, as a source of information. So it's been a bit of a muddle in that sense. I wrote the vision and then I thought well is it actually focussed on the right things? I'm sure I've got some of the big ticket items addressed, but I'm not sure that it's complete, it's not just about being focussed in the right area, it has to be comprehensive as well. So I had addressed the numbers side of things but not the culture side, so the last fortnight has been very muddled for sure.

RD: Can you explain in more detail what you have proposed on the numbers side? What are the analytical functions you were talking about?

CRM: Right, so one of the things that I'm looking to do in the vision is to come up with a hierarchy of risks, an architecture for storing the data, and I started thinking that I need to be able to aggregate the risks up the hierarchy, in terms of causation. When I'd been writing my vision, I had envisaged that I could just use like a fault tree methodology, so I thought I'd sit down and write a spreadsheet to test it and see how it works. But as soon as I wrote the spreadsheet I realised its not that simple actually, you can't simply aggregate things.

RD: You were thinking mathematical aggregation...?

CRM: Yeah. I realised that when you're looking at risk you have many possible sources of failure each with the likelihood that it occurs and the consequences of it occurring, and so to determine the probability you need to know all the possible outcomes. So if there were ten sources of failure then in essence what I get is a distribution of ten risks. So if the parent risk was say a pipe breaking due to a seismic event, then there is a distribution of events, this one could be a one in ten thousand year seismic event, and this one here could simply be one in ten year seismic event you know. So you get that distribution of likelihood versus consequences, and then that flows up to another source of failure, say can't deliver water, and that flows up to another business objective. And so I suddenly realised well, bugger, that screws my probabilities, cos if I express them in terms of return periods or something like that,

something like a dam would be a one in ten thousand year event, but then again I could have health and safety events that occur ten times a year, so I could have values greater than one. So even though I can multiply this consequence by this value here to get a risk score for each risk at the lowest level, I can't sum them across and up the hierarchy because they're not commensurate. So that left me with a major technical problem. Cos I had thought that I could take risks here that were described as class one to five and be able to aggregate up to the parent risk and say well it's a class one to five risk somehow you know. I realised that it's still just discrete analysis, dealing with integers, one number here and one number here, where of course in reality there's a distribution of risks, so a distribution of probability versus consequences. And as long as you're using discrete analysis there are going to be some pretty crude approximations in any aggregation of the risk. To go to any higher level of precision or confidence we would have to go to distributions, so to Monte Carlo simulation. But that's a step-wise change in how you tackle the problem. It's more realistic I guess, but I think there's a lot more preparation required before it's possible, and I'm not sure whether the organisation's ready for it just yet.

RD: Have you looked at what's available on the market? Are there software packages you could use?

CRM: Yeah, it did dawn on me that presumably other people would have had this same difficulty and would have found a way to get around it. So yeah I started looking at some commercial software packages, and particularly looking for the phrase 'aggregation of risk' in their capabilities. I was hoping that they would outline a method, but of course but they won't. I did find a couple of software packages that talk about the fact that they aggregate risk, but I guess what makes me nervous now is the thought that is it a black box? I don't know how they're doing it. And if I don't know how they're doing it, that influences my confidence as an engineer. It worries me a little bit, cos I like transparency and knowing exactly how things have been processed so I can have confidence in what I'm saying.

RD: Yes, ok. Me too.

CRM: So it's been a bit challenging. For example, yesterday I went into a meeting with [*Manager, Internal Audit*], who wants to audit some of the health and safety stuff. Now I think we have about ten to fifteen health and safety risks in the risk register, but they're all over the place. There's a risk that we won't comply with health and safety legislation, there's another risk about pedestrians falling off a pipe bridge, and there's another risk about working in confined spaces. I guess you could fit them into a hierarchy of causation or some categorisation like that but they're so higgilty piggilty, there's only fifteen of them and they're all over the place. So I explained to [*Manager, Internal Audit*] and [*Health and Safety Manager*] that I want to implement a structure into the data. I think [*Health and Safety Manager*] picked it up and

[*Manager, Internal Audit*] understands it as well but he's not interested in the detail. So I spent today looking at the risk descriptions that [*Health and Safety Manager*] has in the register, starting with the parent risk and working out how that could be caused, and working my way down to try and test the hierarchy, but leaving out the numbers. But I realised that you get this enormous explosion in the number of risks at the bottom of the pyramid, very quickly in terms of the number of scenarios that you have to deal with. So for health and safety, for example, by the time you had done the analysis and you had tied legislative compliance at the top down to the hazards that exist in the various work places, you'd have well over 300 risks or something like that.

RD: Just for the one business objective?

CRM: Yes. And if you go to three tiers down you'd end up with an enormous body of risks just for one objective, and so I realised that if you go through the process with the kind of rigour that I want you would end up with thousands of risks and you'd only be a small way down your hierarchy. But one of my overall objectives is to simplify the register, and introducing an additional 300 risks would just go completely against that philosophy. Now I want to get better insight into the data but doing that by a mathematical method would deliver a huge data set that would be too cumbersome to manage. We don't have the integration of systems that would enable that data to be managed automatically, so getting that level of insight to support decision making would come at too large a cost to the business in terms of making sure that the data was good quality at all levels. Also, there's a perception inside the organisation that the number of risks in the register is proportionate to the efficiency of the system, and when I've talked to people one of the primary factors influence's their level of buy-in is the amount of extra work that's imposed upon them by the risk management function. So given that we've got eight hundred risks at the moment and the system is so cumbersome, and people are talking about having just a hundred risks, well, I'm not going to be able to sell this. I need to make some kind of approximations to try to get the numbers down.

RD: You mentioned just now about having good quality data. Can you elaborate? What do you mean by "good quality"?

CRM: Well at the moment our risk management function is inefficient. PWC called it a push system as opposed to a pull system. So we're required to push it to make it work, it doesn't prompt us, it doesn't email us out when tasks are required and things like that, it doesn't pull us along in a way. So we have a very inefficient system at the moment, and one of the largest reasons for that is that we have eight hundred risks in the register that are really poorly described. Also our system's not linked with any other kind of business system to automatically update things. There's no architecture for the storage of data. It's all just on one level, eight hundred risks. Where I want to get to is where you go into the register and at the top, at layer one,



we'll have ten risks. So you can get your head around that. And then at the next layer down we'll have say fifty risks. At the moment you can't interrogate the data like that. To get there we really need to clean up the risk descriptions, so I took the health and safety risks and I started at the top with non-compliance against the legislation. So how could that happen? And I just described the causes and I've ended up, I think, with a three-tier hierarchy of causation. They're just the risk descriptions but they're now descriptions that I feel more comfortable about. They're actually meaningful insightful descriptions. So as far as tackling Watercare's risk framework, part of me's been tempted to just forget about scoring any risks at this stage, and to just go in and identify the risks, get the description done, and that would at least give me some alignment, a hierarchy of causation. And once we get to that point we can think about prioritising them. If we had that hierarchy it might be easy to jump to Monte Carlo analysis, or maybe we would just score the risks the same way that we are at the moment, in case we have to let go of this concept of a mathematical aggregation. But at least I would feel like I was making progress with cleaning up the data. Part of me thinks that if I could just get good descriptions of the risks at each level then I would have made progress cos at the moment in the risk register we have a couple of risks here and a couple of risks there, and I feel there's holes in it because there's no structured way that they've gone about, maybe even gaping holes. But one of the constraints on me, I guess, is that we have a monthly reporting process to Board, and the Board are used to having a number of risks reported to them. If I went through this process then I'd have to ditch the scoring process, which would mean that I couldn't report risks to the Board in any summary manner numerically, which of course would then beg the question: why can't you use the existing method?

RD: You mean you couldn't report an aggregated profile?

CRM: Ah... yeah...mhmm... but it's just occurred to me that that's not actually true. If we went through the process I just described then I'd have a certain number of risks, the risk profile if you like, defined at Board level, and then those would expand to a certain number at executive level, and then again down at senior management. So when I reported risks to the Board I could say that they have twenty Class 2 risks and say four Class 4 risks, or something, and that's their profile at that level. At the next level down that profile might expand out to be two hundred Class 2 and forty Class 4, the concept being that when you're reporting a risk, when you're describing a risk profile, you're describing it at a level in the company. Anyway, that's the way I'm picturing how it should work if our risks are derived from our business objectives. But at the moment we don't differentiate between levels. So we've got non-compliance with health and safety legislation, we've got working in confined spaces, and we've got say failure of oxygen, a worker's breathing equipment. At the moment I report these as three risks, but I'm thinking that's wrong. Really it's one risk at this level, one risk at this level, and one risk at this level. But it just made me realise that I

could still look at the risks in isolation and just sum that whole body of risks and report it as it is at the moment, even if that is a crude method of aggregation.

RD: Yes, moving to your hierarchy wouldn't make the current scoring method unusable. You could still give the Board the numbers they expect to see.

CRM: Yup, even though in my mind those numbers mean absolutely bloody nothing. It's like throwing a handful of darts at a dart board and saying there's your profile, it means nothing. But once I had a data structure in place, and architecture, then I could feel more confident that there weren't holes in the data. And then I could work on changing the scoring method and moving to a more sophisticated means of aggregation, maybe Monte Carlo analysis, something that can represent the distribution.

RD: Sorry to jump in. You commented that you think the current representation of the risk profile is essentially meaningless, and I agree with you on that. Does that suggest something about the Board's level of understanding of the risk information they receive.

CRM: I think it probably reflects their level of understanding of the risk process. Then again, you can't expect any member of the Board to fully understand the details of designing a pipeline, or operating a wastewater treatment plant. But I suppose it's fair comment to say that the Board don't have a good level of understanding, perhaps only have a crude level of understanding, given the fact that they've not queried where those numbers come from. So at the moment we present a risk profile to the Board in a certain form, and now I've come in and had a look at that profile and because I like structure I'm thinking well actually that's really crude, the numbers that are being presented actually don't mean anything. But those numbers should mean something. I mean, that's the concept that a company has a risk profile, that there is a shape or distribution of some kind, which represents an understanding of how your risks are distributed. I guess its getting an understanding of where the business carries most exposure, the greatest risk, so where the Board need to be most nervous, where they should feel most exposed, and where they should be pushing for continual improvement. I think that's what the risk profile should explain. So I'd like to impose an architecture on the analysis so that the numbers we present to the Board would actually mean something. I would like to tie the risk profile back to business objectives, instead of business groups, so that the Board could say, "well alright, our objectives are A, B, C, D, and E, and the business objective where we face the greatest downside risk is this one". So next to each objective there'd be a distribution of risks, a measure of how risky they are. But that's a significant step change from where we are now, and how do I convey this to lay people. I mean, they probably love the current system, a simple number. So yes I guess you could question their knowledge about risk. But I also get the impression that not many of the people out there who are actually practising risk management

are even thinking about risk in the terms we've been talking about today. I don't think there's been much thought go into how you can present risk data in a meaningful manner, and the implications that has on how you collect the data, and how you score the data. There are the standards out there, the risk management standards, which basically give you a representation of how risk management works, so you just jump in and go for it, but I'm finding out more and more that it's not simple. When you read the risk management text books they just talk about a matrix. So they put likelihood versus consequence, and the consequences might be measured simply as low, medium or high, or they might go to the next stage with minor, moderate, tolerable, intolerable, and catastrophic. But when you delve further and further into it you realise that it's a hell of a lot more complex. Consequences are not just measured in one dimension, and consequences can be sequential, so one consequence could trigger another consequence and so on.

RD: So you're saying that the established knowledge base out there is not proving to be very helpful?

CRM: Yeah, at this stage I'd say that the texts are too high level, too simplified. And maybe that's a problem. Maybe I'm trying to work at a level that's tackling too much. I don't know, but definitely the publications that are out there, like the guidance standards and so on, are so simple. They provide a very generic framework, but where you try to apply that generic framework to a real business, in a way that you are actually trying to get value out of that process you realise that it's not nearly as simple as it's represented. Ideally, you need to take the generic framework and mould it to suit your business. But that requires an understanding of your business, it requires you to analyse your business, to dissect it in some way. And where I'm having trouble is dissecting Watercare and understanding what's fundamentally important to the business in the first place.



## Dialogue No. 2 – Quantifying risk profiles

**Source:** Transcript 9 July 2007 (CRM:06); 3,200 words

**Participants:** Corporate Risk Manager (CRM)  
Principal Water Planner (PWP)

---

CRM: What I wanted to do today is to understand your planning needs. I've been putting together a document saying this is where I want to take the risk management framework and I remembered quite early on when we were talking about things that you were saying that you struggle with the risk basis for planning...

PWP: Yup.

CRM: [*PWP: responds in the affirmative throughout this passage*] ...that in some situations you had no basis for replacing the assets and you basically just let them run to failure, or until the cost of failure became so obvious that there was an economic case to replace them. So I'm trying to get an understanding of what we can do to help you there, and with the AMP formulation process, so I guess what I can do is talk to you a little bit about some of the concepts that I'm thinking about and see whether or not those would help. I have a Project Improve presentation coming up and one of the things that I'm going to propose is that we actually investigate and try to define what I would call risk acceptance criteria. So at the moment we have risks that are classified one to five and there's a perception that class four and five risks are unacceptable. Now I'm not necessarily in agreement with that. I actually think that all we have got there is a prioritised list of risks from low to high and what I would like to do is actually go down and be able to say lets get the Board and all our other stakeholders to sign off on a line single line through our body of risks that says this side are acceptable and this side are unacceptable. Now I don't know where that's going to sit in terms of our current classifications but what I'm hoping is that by giving you a definite line it will somehow improve your business process. I'm also trying to work with the condition monitoring people to tie their work in, and what I see happening is that we go out and do asset condition monitoring and then that information somehow, whether intuitively or through analysis, would give us a likelihood of failure. Then if we know the consequences of failure we have a risk, obviously, and the question then is whether that risk is acceptable or not, and should we therefore replace the asset. Now I know that we do this at the moment and we kind of get to some judgement here, but I don't see that we've got any kind of solid criteria for making that judgement, particularly for smaller consequence events. I know that when the consequences get right up there like the Hunua No. 3

they're quite happy to say that's unacceptable but in terms of our more day to day refurbishment or replacement of assets I don't know that it's clear.

PWP: Well it's really a consensus process between Operations and Planning because I don't always agree with their outcomes.

CRM: From the condition monitoring?

PWP: Yeah. So they assess the pipeline or whatever and see that it's maybe knackered in a few places, or that it could be knackered in a few years, and you know engineers are sort of emotionally attached to their work and always want things to be replaced, but it hasn't failed, hasn't given us any hassles up to now, hasn't cost us any money, and we can't just bring it forward cos it's a ten million dollar project...

CRM: Can't just pull ten million dollars out of the air...

PWP: Yeah, cos the budget eventually has an effect on the water price which is a political issue, so we've almost got to protect the AMP from our point of view.

CRM: So you've only got a fixed pool of cash to play with?

PWP: Almost yes. We've got to make sure that the projects that go in the AMP, especially in the first five to ten years, actually need to be there. So often they will do the condition assessment and say "this main is critical, we've got to replace it", but when you look at the main in the context of the whole system it's a non-critical main. So from our perspective then let's postpone it, you know, until after it's failed a few times and you actually have to, until it's giving you too much hassle. An example is the Upper Ngaio main, an early 1900's cast iron main which the asset condition specialists have said needs to be replaced. Ok so it probably needs to be replaced cos it's more than a hundred years old, but it's still functioning ok. It's really only when you operate the system a little bit harshly, and you open and close the valves too quickly and create transient pressures that it may fail under those conditions. But if you operate the system well then it's not a problem, and anyway even if the main fails it's non critical cos there are three other dams feeding Huia treatment plant.

CRM: Right ok. And when you talk about criticality, in terms of a pipe in the network, how do you assess that?

PWP: Well criticality is related to the consequences. Well actually it must be related to both consequence and likelihood. I mean then higher the risk the more critical. So it's not only the consequence. If you have a main that's failing every month, or even twice a year, then that becomes critical because...

CRM: Because of the costs it's chewing up?

PWP: Yeah because of the costs and the impacts on customers. Even if it's just a small number of customers, just a few thousand people at a time, but if it occurs all the

time then obviously that becomes a fairly critical issue. So it comes down to the risk score. And what the asset specialists are saying is that the likelihood of failure is high, but we have to also look at the consequences in the context of the total system and assess whether the consequences are low or high, plot it on the risk graph, and then decide whether we need to replace it or not.

CRM: Yeah ok, alright. So in theory I'm assuming that when the guys go out and they do the asset condition monitoring, that they should be able to plot age versus condition for different types of assets and get some kind of crude correlation about how we expect those assets to deteriorate. And then over time with good data I'm assuming that we could interpret that correlation into an age versus likelihood score so that we could predict the likelihood of asset failure a couple of years in advance with a higher degree of certainty than what we have at the moment. Well, no sorry that might not be true, but at least with a more analytical basis.

PWP: Yeah

CRM: And I'm assuming then that if we can predict, based on asset age and type, roughly when we expect it to reach a certain likelihood of failure, and if you know the consequences of that asset failing, then you would be able to work out when it was going to pass over some predefined risk acceptance level. Now I don't know whether that's, I mean, to be honest, I'm just trying to provide analytical data for what we currently do anyway.

PWP: Which will be great, I mean it would be very good, yeah.

CRM: Ok well that's the fundamental question isn't it, is it actually going to help?

PWP: Well, it's a question of calibrating the processes against the analytical data, calibrating the outcomes that we have now against the analytical data and then improving that calibration all the time.

CRM: Ok good, cos at the moment my perception is that even though you've got the asset condition and what risk that represents, you also kind of have a debate about whether that risk is acceptable, about whether if it's a class three risk you should replace it as opposed to, well I don't know, but maybe having a more defined threshold would help. So how do you use risk, or rather how would you like to be able to use risk data in your AMP formulation process?

PWP: Well the AMP covers twenty years and normally we only have a snap shot of risk at the current date. It's actually hard to predict risk into the future because you have to develop those snap shots almost at each year going forward and try to predict how the risk will increase for different assets, and at what point in time we're going to exceed the risk threshold. So its fairly easy to do risk from a growth point of view because you know you have to provide a certain service level at some point in time and so basically as soon as we reach capacity we have to upgrade or replace the asset or duplicate it. So from a growth point of view what we're saying is that if we

don't do that in time then the risk of not being able to supply to the required service level is so high that we can't tolerate it.

CRM: Yeah that's right yeah.

PWP: But from a condition point of view it's what you've drawn there, some relationship between age and risk needs to be drawn up for each asset so that we can predict when the risk will become intolerable.

CRM: Ok, and from what I understand you've got a certain pot of money and there's a desire to be able to juggle the projects so that you get the maximum overall risk reduction.

PWP: Yes.

CRM: Right, so to be able to predict the risk reduction versus the amount that we'd pay for that, and to maximise or optimise the impact of that.

PWP: Yes, that would be ideal. But getting the information in that would allow you to do that, that's the problem.

CRM: Yeah the quality of the information, the data quality's not good at the moment.

PWP: Well the AvSim model that I mentioned to you the other day already has typical relationships built in, the Weibull curves, for different types of assets. So you can start off with those and then calibrate the curves for individual assets as you go into the future. The problem is that our mains don't fail a lot so there isn't any information to calibrate it against unless you do specific condition assessments.

CRM: Right ok. Now I was wondering about that in terms of formulating the AMP. So if you know your current risk, and if you're implementing a new control then you should get a lower risk in the future and that then gives you the change in the risk profile, the delta risk for a project. And we can put that against the spend. But I was thinking about how to predict the future risk levels using the condition monitoring data like I talked about earlier. I was thinking that the only way you could do that would be to model the network in the future. Currently we measure our consequences on a five different scales, but they're quite subjective, particularly the last one, Asset Management, and in order to be able to model consequences we'd need to get rid of some of that subjectivity. So I've been wondering whether or not we can correlate the consequences of asset failure with the level of reliance that we have on an asset in the network. So for example if you've got a network model solution which says 70% of Auckland's water is going down the Hunua No. 3 in 2006 and in 2020 is going to be 90%, then you could somehow use that to measure your consequences of asset failure. I haven't really thought about it in much depth, but I'm assuming that we can do that...

PWP: Well what they've done with the AvSim model is to create a proxy for the effects of failure by using the network model first. So what they did, and I haven't worked



through it in detail myself, but they looked at how much water is running through each of the network pipes, so the network model represents the process flow, and then if you break a pipe what are the resulting influences on the network.

CRM: The network model, that's a hydraulic model is it?

PWP: A hydraulic model yeah. So you can go "lets break a pipe here" and then see what the effects are downstream. You can run an algorithm through the network model and basically get an output of the effects of each individual pipe break, so that tells you basically so many people would be out of water as the consequence of any particular pipe break. So then those results have been built into the AvSim model and translated into a consequence score.

CRM: Sorry, so you take the output of the hydraulic model and you feed that result into the AvSim model?

PWP: Yup. Basically it's just the volume of water that will be cut off when you break a specific pipe. That translates into a population number, so how many people would be affected, and that number gets fed into the AvSim model, and then translated into a consequence score through the risk framework.

CRM: Right ok.

PWP: So what you have then is a process model where you can take any individual asset out and see what the resulting consequence score would be. And the AvSim model has the Weibull curves built in so that provides you with a likelihood of failure for that asset, so you can use the model to assess the criticality of each individual asset. Provided you have reliable consequence information and reliable condition information you can get a prioritised list of assets by criticality.

CRM: Ok. I've got a meeting with [*RCM Manager*] to understand the RCM and AvSim stuff better so I'll talk with him about that. But I guess one of the things that worries me is, say you've got twenty possible projects for the AMP and you can't fund them all so you've got to cut ten of them or something like that. If you're modelling the future network configuration, with twenty possible projects that's a hell of a lot of different combinations for what the future network could look like and to assess the future reliance on an asset depends on all the projects that could have been done up to that point in time. So I'm wondering then how easy it is to do this, to adjust the model and rerun it. And I had thought we might be looking at running either the hydraulic model or the AvSim model, but I didn't realise that one fed into the other. So I'm assuming that's a hell of a lot of work to change network configurations, and re-run all the models, and I'm wondering whether it's practical.

PWP: Ah, yeah there would be a little bit of work going into that. I mean we've already got a network Master Plan that's staged in five year increments over a twenty year period, so we could basically create an AvSim model for each of those. But it's an iterative process again...

CRM: Very much so.

PWP: Because you'll run the AvSim model, get the criticality out of the AvSim model, and then you see which pipelines need to be upgraded, put those back into the network model, and then the network model gives you a different result.

CRM: Right so its changing, that's right.

PWP: Yes, the characteristics of the network are changing.

CRM: Right so five years later those characteristics have changed quite a lot.

PWP: Yeah but you know that's thinking about it in a very complicated way, I don't think we'll go to that extent.

CRM: Do you think we could do it in a cruder way?

PWP: I think we should just do it for the current network and for the current AvSim model and that will probably give us enough information to populate the AMP. Then the key would be to get the condition information, because the whole purpose of the exercise would be to use the Weibull curves, so if you follow the curve going into the future it will tell you when the likelihood of failure increases so much that you need to look at replacing the main.

CRM: Yes exactly...

PWP: So to use the curves to predict when a pipeline needs to be replaced. You can use just the default curves, but then you'll get to that point in time, ten or twenty years in the future, when the model says the pipeline should be replaced and it hasn't failed...

CRM: So you go and do a condition assessment, and you can then say "well actually we could get another five years out of it".

PWP: Exactly, yeah.

CRM: Ok. Cos the thing that has worried me about going down this path is that I can see it getting incredibly complicated. I guess what I'm wondering is whether we can take a cruder approach to it, you know if it's too bloody complicated to do every year or even every five years then could we take a step back and do it at a cruder level then we could still have a degree of comfort that we were in the right ball park.

PWP: Yeah. I think you can do it for the first year and just work through the process and see what we get out of it. Then we can decide whether we need to make it more complex.

CRM: Ok, alright. I think it's definitely worth having a look at because, and I need to learn more about RCM and AvSim, but I see those models as being the repositories for asset condition information within the company. So we don't need to have all that data in the corporate risk register, in fact we will never get that level of resolution in

the risk register, but it is possible in those models. So we can use those models to provide the detail or resolution for key risks in the register.

PWP: Yeah, well eventually if you have the whole water system and wastewater system built into AvSim then essentially you'll have a prioritised listing of assets numbered from one to thirty thousand in terms of asset criticality.

CRM: Yeah and I don't want that level of resolution in our risk register. I think that's just going to kill us but if it's in a tool like RCM we should be able to tap into it. But to do that I need to understand exactly what's going on with the models. I know that Paul understands it, but I'm not sure that anyone else in the business really understands it. It's really a bit of black box magic at the moment. Alright, well that's good. I really just wanted to check that what I'd been speaking with other people about and that what I'd been thinking was actually along the right lines. It's still very high level still but...

PWP: Well as far as I'm concerned it sounds like you're on the right track.



## Dialogue No. 3 – Proposed developments

**Source:** Transcript 20 July 2007 (CRM:09); 2,800 words

**Participants:** Corporate Risk Manager (CRM)  
Richard Donnelly (RD) [*Researcher*]

---

CRM: ... so I talked to the project managers and they were having trouble measuring project risk because they were trying to use the parameters that are of interest to the Operations guys to measure risk on a project. They were saying they'd like to measure risk with the parameters that they think about on a day-to-day basis... So I ended up proposing the project risk register that measured risk in five parameters that were relevant to projects, and then coming up with a method to convert that to how those risks would be significant at the corporate level. Now I was trying to work out why that's a step forward cos they had a project risk register before and I've just given them a different one. Then I realised that what I'd done was to provide them with a tool that's appropriate for their business focus, that's right for their context. So in the past they had a tool that was designed for another context, the Operation's context, and of course it wasn't giving them sensible results. It was very difficult for them to apply it so they had to manipulate the numbers or they'd built in a fudge factor so that what they ended up getting out of the tool reflected what they intuitively felt. Now what I've done is I've provided them with a tool that is using the parameters that are of interest to them, it's appropriate to their context, so they shouldn't need to manipulate the results. It should give them insight into the issues that they need to worry about rather than them having to use their intuition to correct the data. I realised that's where the value lay. Giving them a tool that's appropriate for their work should raise the level of assurance in the data because there's less manipulation or fudging.

RD: By assurance you mean the confidence that decision makers can have in the risk data?

CRM: Yes I think so. Because the project managers think in terms of time, cost and quality all day, by asking them to quantify risks in terms of time, cost and quality, they're familiar with them. So the risk measures in the project risk register should reflect their intuitive feeling, without them needing to apply fudge factors and things to distort the resulting risk scores. And I guess that emphasises the importance of context, that you should be deriving your risk management frameworks from your context, which is what the [Risk Management] Standard says. So I'm going to recommend that the company moves away from a one-size-fits all approach to risk assessment, we need to differentiate between the Operations context, the Project

context, and the Strategic context, and I'm going to use the project risk register development to illustrate that.

RD: Now that's interesting, looking at your approach there. You've been talking for a while about having a hierarchy of risks, about constructing risk frameworks from a deconstruction of the corporate objectives, but, what you're saying with the project framework is that it's necessary to understand the specific decision making needs in the different operational contexts, and then somehow link back up to the corporate objectives.

CRM: Yes, ironically, I had go right down into the detail of the operational context in order to understand something that is really the very first step of the risk management process in the Standard. I think actually that there's a need to understand both ends of the spectrum, you can't just start purely at either end. A good example of that was, the other day, I had been looking at the corporate objective for health and safety, which was that there should be no successful health and safety prosecutions. I was trying to deconstruct that and I was thinking, how could the company be prosecuted? Perhaps an inspector deems that you have an unacceptable level of safety, or there is an accident which triggers action. And then I thought, well ok, who could get injured? It could be the public, it could be an employee, it could be a visitor. So, how could someone be injured? Well, they could be exposed to hazards. And, initially, I was thinking do I have to include this in my risk register, cos that could be enormous; there must be hundreds of health and safety hazards in an operational environment like Watercare's. But then I sat down with [the Health and Safety Manager] and he showed me the hazard register that he maintains for compliance purposes. He described how he goes through a process of identifying where we have health and safety hazards, what work processes and locations they arise in, and the associated assets, and then he assesses the severity of them. So this hazard register, which is a spreadsheet, contains about three hundred odd hazard evaluations, and I realised that it's just a form of risk register. In essence, each hazard evaluation is a risk assessment. Bingo! All of a sudden I've got a link there to an evaluation of all the health and safety hazards in the company. My corporate risk register doesn't need to duplicate the existing hazard register. I could just go down to a certain level and then point to or reference the existing tool. Then it dawned on me that I can have a corporate risk register that sits up above a whole lot of operational tools; so, for example, there's the health and safety hazard register, there's also the public health risk management plans, and the Reliability Centred Maintenance models. So the concept is that I take my corporate risk register down to a certain level and then reference tools that are already used and maintained by the different business units. I think the value there is that I'm making a link between corporate objectives and tools that already exist, to knowledge that already exists and is managed on a day to day basis by people in the company. I don't think there's any point in trying to draw that information, that level of detail into

the corporate risk register because that would just be doubling up. So that's a further development, but again, if you look at that, it's come out of looking at the needs at an operational level. It wasn't until I looked at the operational context that I was able to marry it together with the enterprise objectives. My approach actually stands out like a sore thumb there. What I've actually done is gone to the operational level and said 'what are your needs?', 'let me understand your business', and while they're talking I'm trying to think of their needs in terms of a risk context. I guess you could say I've very much tried to go along the bottom and say 'what are your needs?', to identify those and look for commonalities and themes, and then asking what does the risk management function have to deliver in order to assist with meeting those needs, instead of coming in and saying 'enterprise risk management is this', and imposing it on the business.

RD: Why do you think you took that approach?

CRM: I don't know. Perhaps, from a personal point of view, I've always heard that historically the guys who were really effective at running companies were the guys who could walk down to the workers on the shop floor and ask 'what do you do here, what are your needs, what stops you doing what you need to do?' Which is what I've tried to do, and I guess my job then is to go away and try to link that with the corporate objectives. So I think the value that I add is in understanding what they need at the coalface in relation to the corporate objectives, and being able to put it all together in an efficient and consistent structure, which I think is the essence of the theory of risk management. It's about having an objective way of measuring things across an organisation, full breadth and full depth. Actually, I have another good example. If you remember, I was talking to [members of the planning department] about their business needs...

RD: Yes, I remember the meeting...

CRM: ... well, we were talking about an analysis tool, a tool to help them formulate the [Asset Management Plan]. It's essentially a tool that would let them look across all the projects on the books and formulate, for different scenarios, combinations of projects, the projected risk reduction for a given level of expenditure. At the moment, that's very much a qualitative process, and they want a tool to provide greater detail of analysis. I realised that, at the moment, they haven't got sufficient quality risk data to be able to conduct that kind of macro-level analysis. And I've been thinking about how to clean up the data quality and I'm starting to see that the data might actually exist elsewhere in the business at the level of detail that they need. With the AMP formulation tool, the key challenge is to predict future risk levels, which relies on asset condition information, condition monitoring and some numerical modelling. That's the same data that the [Reliability Centred Maintenance] models need to have. So I'm thinking if those models need to have data stored in a consistent way, and if they store it in the right way then they'll satisfy a need elsewhere in the

business. So maybe the necessary data doesn't exist in the risk register, but then maybe it doesn't need to. Maybe I just need to be able to link the different areas of the business together, or maybe my role is identify where a need in one part of the business could be addressed by using or modifying a capability in another part of the business.

RD: Hmm, looking for the gaps no one else is seeing.

CRM: Yes, it seems so. A related development there is I'm looking to introduce a decision uncertainty analysis on large [capital expenditure] decisions. Now this isn't something that the business knows it needs. This is something that I saw at a conference that I went to and I think it would be a valuable function. It's basically about calculating the percentage reliability of your cost-benefit estimate. So rather than having just one number, you've actually got a distribution of potential answers given the uncertainty of the parameters that you use, and you can produce a cumulative probability distribution of outcomes. So, when it comes to big [capital expenditure] evaluations I think we should be having a presentation of the uncertainty associated with those evaluations. At the moment, [the CEO] or the Board get presented with a single figure for the [net present value] of the decision, and I'm going to advocate that we use some Monte Carlo simulation work to produce a distribution of possible outcomes given the uncertainty that we think exists in the underlying parameters. Then we can show the confidence that we have in the project results. I'm not anticipating that this one will be hard to sell to management. Similarly, another development, which might prove a harder sell, and actually this is something you have previously commented on, is the use of the root-mean-square averaging method to produce a single risk score from five different consequence scores.

RD: Ah yes, I had suggested that it wasn't an appropriate method because it obscured the variability in the underlying consequence scores.

CRM: Exactly. The RMS average is a nice way of consolidating a lot of information. It produces a single number which is easy for the managers to focus on, but at the same time you lose a certain precision, or depth of knowledge, and there could be situations where that is dangerous. Quite often you end up with risks that are very much one-dimensional or heavily biased in one or two consequence categories.

RD: Yes, so maybe the primary consequence is a health and safety impact. Someone could get killed or seriously injured, but the severity of consequences in other categories is minimal, and the RMS average doesn't reflect that one-dimensional severity.

CRM: Yes, exactly. Well, I want to introduce risk tolerance criteria, to make it more explicit about when a risk is and is not tolerable. But that's really impossible for a parameter which is an average of five different types of consequence. I think it is necessary to



define risk tolerance criteria for each consequence category, because you could conceivably have quite a different risk tolerance profile for financial impacts than for, say, health and safety or service delivery impacts. That means you need to technically derive a risk tolerance definition for each different measure of consequences so that, when it comes to evaluating your risk, you might have a risk that passes on five out of the six tolerance criteria but fails on one. So, maybe the average of all the consequences is relatively small, say a Class 2 risk, but unfortunately it has a peak in health and safety, so maybe it's unacceptable in the health and safety context. I guess that makes the analysis a bit more complicated but I think it's all doable, it's just mathematics and data manipulation. The difficult part will be how to represent and communicate risk in this form. I mean, rather than one number for each risk in the register, we'll now have five or even six depending on how many consequence categories we have. With the new method we'll have six different plots of consequence versus likelihood, and each one will have a different distribution in that dimension. Now that's a big increase in information and I need to find a smart way of communicating what has been an expansion in information into something that is visually simple, that the directors can quickly scan and identify the key priorities...

... So those are the main functional developments that I'm going to propose in the presentation next month. I guess what is foremost in my mind at the moment is how to take all these disparate ideas and put them into a cohesive presentation where each idea in itself illustrates practical value but also that, as a whole, the set illustrates that I'm worth listening to, that what I'm proposing is suitable or appropriate for the company, and not just a bunch of academic wank.

RD: And, do you have any thoughts on how you're going to demonstrate that value?

CRM: Well, it's actually very hard. I mean, I've consulted with various people in the operational groups to find out what their needs are, and I think I've got a good idea now, or at least a better idea, of what their needs are, and how to provide for those needs analytically. In some cases I've got very good definition about how to do it, but in other cases, while I think it can be done, in fact I'm very confident it can be done, I'm just not sure yet exactly how to do it.

RD: When you say 'analytically' and 'you've got good definition', you mean...?

CRM: Ok, so I think that if I sat down I could write out the mathematical process for delivering the result. I have an understanding of what data is required, and the process of how to manipulate that data to get what the end user requires. In other situations I think I know what the end user requires, and I'm pretty sure I know the input data, but I'm just not sure how to process it. So, in those cases, there's some uncertainty about how to deliver the outcome that the end user wants, although I feel confident that it's possible. In terms of my presentation, I think that there's probably about five or six individual developments that I feel I could talk about the value of,

but where I'm struggling is with that one slide that says these are the three things that I'm focussing on and those are the three things that give management the confidence that I've got my eye on the right things. I mean, I have a series of what you could call operational tools which I think individually will add value, but I'm struggling to wrap them up as a cohesive set, to be able to say these actually represent three fundamental improvements leading to improved decision making or something like that. It would be nice to be able to stand up and say that over the next twelve months these are the three characteristics I'm going to focus on, and dropping out of this work will be these six advances in the business functionality, but I haven't got that clarity in terms of tying it together yet.

## Dialogue No. 4 – Engaging with stakeholders: Operations Managers

**Source:** Transcript 29 August 2007 (CRM:11)

**Participants:** Corporate Risk Manager (CRM)                      Wastewater Treatment Plant  
Water Networks Manager (WNM)                      Manager (WWM)  
Water Treatment Manager (WTM)                      Environmental Services  
Risk Specialist (RS)                                      Manager (ESM)  
Laboratory Manager (LM)

---

CRM: There's a whole lot of development work that is going on and will continue for the next couple years on the risk management framework. What I'm looking to do today is to outline some of that development work that will have an impact on you guys, and also to get your feedback on some of the development that has been done to date. Have you guys had a chance to read the document I sent out?

[general responses]

CRM: Ok alright. So the risk management steering committee have signed off and endorsed a programme of development that is centred on improving four areas of our risk management, the first of which concerns the frameworks that we use to quantify risk. The other three which follow on from that will target improvements in the quality of our risk data, which will then help us to improve our risk analysis. But what I'm talking to you about today is that first area of development which is improving our risk frameworks. At the moment we've got one risk scoring methodology that we use across the whole business and if you've used it you would have discovered that in certain situations it is very subjective. This has made it difficult for certain parts of the business to use the framework and get a reasonable representation of their risks. And that is having a big impact on the resulting data quality, and in turn that has big impacts down stream in terms of what that data is used for. What we're looking to do is to move away from a one size fits all framework and to put in place risk scoring methods for different areas of the business that better suit their needs. We're proposing that there will actually be three frameworks. There will be a risk framework designed to suit what we're calling Business Operations. There will be a risk framework that's designed to suit what we're calling Business Projects, where by "project" we mean anything that is developing or changing the organisation's capabilities. And then finally there will be a framework at the very top of the organisation looking at strategic risks. And the way that risk is defined and measured will be different for each of those frameworks. The idea is to provide you guys in the business operations area with a framework that consists of

the kind of things that you're thinking about on a day to day basis, so when you're asked to assess risks you'll be able to use parameters that you're familiar with or concerned with on a day to day basis, rather than having to think about parameters that are foreign. So if I look at our current framework for a moment, we currently measure consequences in terms of reputation, finance, environment and public health, health and safety, and a category called asset management. Now asset management, from a water and wastewater service point of view, is probably the most important parameter because it's the one that is used to capture the impact of service interruptions. However, it's also the parameter that is the hardest one to quantify, and this makes the existing criteria very subjective. And it seems to me that measuring consequences in terms of Reputation double counts incidents that we've already captured. So if we have a break in a major pipe and the public are exposed to a water supply interruption, the reputation category seems to double count that impact by how it is reported in the media. So you can question whether Reputation is an appropriate way to measure risk. So if you look at the system we have at the moment you find there are quite a few limitations that are having a big impact on the data quality, and that has been the motivation for reviewing the framework. What I set out to do was to craft a framework that would provide staff with parameters to measure risks that they were very familiar with. Now if someone asks "what are your risks?", then your first question should be "risk to what?" And it is always about risks to achieving something. This means that in terms of putting together a framework for a business the risks should all be aligned with what that organisation is trying to achieve. For Watercare, there are eighteen performance objectives specified under the Sustainability Policies in our [Statement of Corporate Intent]. And I believe they can all be condensed under five basic parameters for business operations, which really are the five things that everyone talks about as being Watercare's core business, and on that basis then, we should be measuring our risks against them. Now those parameters are the ability to supply water, to provide water of an acceptable quality, to maintain statutory compliance, which is where I want to get feedback today in terms of our wastewater service because from what I can see the primary objective there is really to comply with resource consents. The other thing that is consistently throughout the SCI is that we have to be an efficient business, and then the fifth one is about maintaining the confidence of our shareholders and customers, which I guess is similar to the current reputation category. So those are the parameters that I'm proposing we should be using for measuring risks, and I think they should be more in line with the kinds of things that you're thinking about on a day to day basis in the Operations area. But before I can set these in stone I need to get your feedback because you will obviously be inheriting this framework.

WNM: Just on the water side, to kick it off. What do you mean by water palatability? Do you mean taste and odour?

CRM: I think so yeah.

WNM: Well if we start with the “Negligible” category down the bottom, I think the key thing is that anything that we have to report to the [Ministry of Health] is a transgression, it’s a non compliance. If we have a non compliance, its not negligible, it can’t be negligible.

WPM1: Well strictly speaking if it’ll cause a downgrade, or incur greater than 10 demerit points, for you that’s a B grade, but for me it’s a yes/no. We can’t get a B in Treatment, we can go from A to C to D.

WNM: Exactly, if we incur between five and ten demerit points that means a drop from an A to a B grade, and maintaining an A grade is the first thing on the SCI. I would have thought that was...

LM: ...catastrophic...

WPM1: Well, is it? That raises the question, is losing you’re ‘A’ grade a catastrophic failure?

CRM: No no, yes. Well that’s exactly the point. Losing your ‘A’ grade is different from poisoning 50,000 people.

WNM: It depends what you lose your ‘A’ grade for...

WPM1: Well...

CRM: These are the kinds of things that we have to balance. So the first question to ask you then in relation to water quality is: what are the primary things that we think about? One is obviously grading and compliance, and the other one is our potential to poison people?

WPM1: Not poison. I think there is potential to make people become unwell.

CRM: Yes.

WNM: But if you compare say under Financial, unbudgeted costs of \$1000 dollars as negligible versus giving 50,000 people shitty water. They just don’t compare. They’re not even on the same scale.

WPM1: 50,000 people having shitty water will make television news.

WNM: 50,000 people getting shitty water will... well [WTM] or I probably won’t be around.

CRM: These are taste and odour issues?

WPM1: Well take the dirty water event on Christmas Eve in downtown. That was less than 50,000 people...

WNM: For the last taste and odour event we had, there were 123 complaints from consumers and Metrowater were screaming. They wanted me on my knees for 123 complaints.

CRM: Ok. Well all the figures are open to flexibility, but that's the kind of feedback I want to get. But it's more important to me at this stage to know that I'm providing you with the right indicators, are they the right things...?

WPM1: Well let's take the grading issue again. Is it really relevant if there's people getting sick? If people don't get sick, but we've put out Boil Water Notices across the whole city to prevent people being sick cos that's our contingency measure, well then the company will change. It won't be WSL as we know it...

WNM: No, exactly...

WPM1: ...the senior management will be cleared out. It doesn't matter... if people get sick that just makes it even worse.

WNM: Yes, the brand will be irreparably damaged.

WPM1: Well, Sydney Water was when it happened to them.

CRM: Yes, but presumably that's a hell of a lot more serious than having a downgrading...

WPM1: No but often you'll get the downgrade anyway....

WNM: The downgrade will be as a result of that...

WPM1: ...because the downgrade is a measure of risk. To get an A grade...

CRM: Granted. I understand that, that you do one you get the other, but not in all situations right? You can get a downgrade in points...

WPM1: Not in water treatment.

CRM: Oh really?

WPM1: Water treatment is yes/no.

CRM: Ok, so the point being that we might have to treat water treatment and reticulation grading separately in this scale.

WPM1: Mhm, they're completely different.

WNM: I think the grading is where you're double counting. I'd get rid of the grading in here because grading is as a result of failing drinking water standards or failing some other measure like, for instance, you could have unqualified staff or you could lack adequate management systems, or you might not have an approved public health risk management plan...

CRM: Right, ok. So what you guys are telling me is that in the water quality scale here, you just want six levels of non-compliance...

WPM1: Well...

WNM: Ah well, no, cos there's aesthetic things. Something that's negligible or moderate can not be a non-compliance. As soon as you get non-compliance it's an issue. Like, for instance, the bacto and protozoa water quality breach that you've got on

the scale there for less than 2000 people and the next level is 10,000... well we can never affect only 2000 people...

WPM1: What about Onehunga?

WNM: Our smallest bulk supply point is 2000 people.

CRM: Right ok.

WNM: So I think there is a difference between a transgression of health significance, and then perhaps a transgression of an aesthetic guideline, so dirty water. Taste and odour is a contractual requirement, it's a guideline value.

CRM: Ok, so in terms of measuring risks for water quality, how do you want the potential consequences described? Cos on the one hand you're talking about compliance with drinking water standards, or the bulk water agreement, and the other one that we have used historically is the number of people that could be exposed. We need to agree on indicators and a scale that you guys are going to feel comfortable with for measuring risk, that will cover the full spectrum of consequences, at least under water quality. I mean, some of what I'm hearing is that it's all compliance issues, it's just the severity with which you non-comply. But I wonder whether you will get situations where you have non-compliance occurring but where it actually has no consequences downstream in terms of the public health. Perhaps because we can repair it quickly.

WPM1: Well I mean the ultimate worst case scenario that we could have would be non-compliance with DWS associated with widespread illness across the Auckland region. Then I think we would be facing prosecutions...

RS: Yeah if it got to that stage someone might look to make prosecutions.

CRM: Yeah but I'm not worried about the follow on consequences. Just the water quality breach, what's the worst that could happen down to the most negligible?

WPM1: Well I suppose the worst would be people getting sick. Then, if you want it in varying degrees, did we have non-compliance with drinking water standards. After that would be non-compliance with drinking water standards where we had to issue boil water notices or put in place restrictions or what ever, but no one got sick. After that it would be non-compliance with one part of the DWS...

CRM: Because, am I right, there are situations in the standards that are more associated with the reliability of providing water quality rather than actually breaches in quality, thinks like the quality of staff...

WNM: Well that's a grading issue. But, for instance, you could have a breach of a THM which is an organic disinfection byproduct and you could have a breach of that standard, but no-one's going to get sick because the impacts of that are measured on a life-time exposure. People aren't going to get sick from a single exposure,

versus you could have a breakthrough of e-coli at Ardmore and a whole lot of people could sick instantly.

CRM: Or this issue of qualified staff, if you were determined not to have qualified staff would you not get grading penalties, even though it has no impact on public health...

WPM1: Well let's face it, under normal conditions at Ardmore you could probably get away with it. But it's during the event, during the bad water quality, when you've had systematic failures, those are the days you need the qualified staff. So you've had people in charge that can do the job nine days out ten, but on that tenth day they just can't. And that's what you typically find has happened during all the water quality events that have occurred in the world. And that's what we're finding too. If you look at the events that we've had at Ardmore and Huia, up until recently it was the same people going out and making sure we got through the fire, cos we left it to other people that we thought were competent but turns out they weren't.

CRM: Ok alright. Well what about water quantity, in terms of the delivery of water?

WNM: Well, starting at the bottom, interruption of supply, so that's no water to 10,000, a 100,000 population hours, that's 10,000 people with no water for 10 hours, and that's negligible risk? That's not going to fly...

CRM: Again the numbers are flexible, but using that parameter of population hours?

WNM: I think it's a good concept...

WPM1: It's easier to define...

CRM: Ok. What about the concept of repair time? Say we have a break in the network that's going to take 12 hours to repair so that could be potentially 12 hours interruption to supply except that we have local reservoirs so you might be able to get 4 hours of supply out of those. Is it feasible to say, given a repair time of this, we have enough buffer in the system to reduce the interruption...

WNM: That's where you'd put restrictions in place. And before you get to restrictions there is reduced supply...

WPM1: I believe reduced supply to be moderate, it's survivable. I mean if the Hunua No. 1 burst had been really bad and we had to reduce supply to certain areas, that's far less damaging to the company than if we ended up with no water. Cos once you get to the no-water situation then you start reverting back to the water quality issue...

WNM: The consequences of the Hunua No. 1 were by definition very serious, unbudgeted costs of between one and ten million bucks, and the cost was a four million dollar exposure...

CRM: Yes, but in terms of water supply, you're quite comfortable with population-hours as a concept, and maybe we need to look at including indicators for...



WNM: Well there's three things: no water, low pressure, restrictions...

WPM1: What about dirty water?

WNM: Well that needs to be under water quality.

WPM1: What about dirty water caused by no water?

WNM: Well that's where all these things interrelate...

CRM: Well then it'll turn up in both columns.

WNM: With water quality, the bottom three, negligible, moderate and important are all aesthetic determinands. The top ones are all non-compliances and breaches.

WPM1: What does a dam failure come under?

CRM: Well under compliance issues there would be statutory compliance, there would be multiple fatalities...

RS: ...A...

WPM1: Well we know it's 'A', but which one?

ESM: Right up the top there catastrophic and prominent breach of statute, I don't know what the statute would be...

CRM: You've got multiple fatalities, and you've got a huge financial cost...

WPM1: Well there might not be any fatalities...

CRM: Well multiple exposures if you like, and the financial loss will be huge.

WPM1: They're almost outputs, they fall out of what's occurred...

CRM: Yes that's right, they do.

WPM1: In some ways the last column is the third floor risk profile. It's what the people upstairs use to quantify risk, where as we're trying to apply a logic to work out what the consequences will be...

CRM: Well that column will always be subjective. But [WNM], you were telling me about the Hunua No. 3 valve failure. There was no real impact on the public because you got it all done in time, but there was an impact in terms of stakeholder confidence, in this case our customers rather than the public, because they started wondering how exposed they could potentially have been.

WNM: Yes, it's the same with the Hunua No. 1 failure. That could easily have caused a draindown if it weren't for the responsiveness of the people who were there.

WPM1: Well if it had happened at seven o'clock in the morning it would have, because you wouldn't have got there in time.

WNM: Yes you're right. So there's a timing thing, and also the people who did respond. So one of the things that's important when you actually assess these risks, you don't know who you're going to have available, or when the event might occur...

CRM: And that actually comes down to how you define the risk. At the moment we are not very good at defining risks. We often just say "the Hunua No 1 breaks", which is not actually a definition of risk. What you're supposed to say is these are consequences caused by some event. So, for example, a failure of the Hunua No. 1 causes a draindown of the network because it fails at 7 o'clock in the morning. That's a risk. It requires us to be more detailed. And these are the problems that you guys are having when you sit down to talk about risk data... Alright... [writing]... I think that's enough on water. Let's skip business efficiency and look at statutory compliance. Earl, when I've looked at wastewater, both for the networks and treatment, it seems to me that if something goes wrong it all boils down to either a breach in some resource consent conditions or some impact on business efficiency.

WWM: Well it will impact on people.

CRM: In terms of public health, is that what you mean?

WWM: Well it could be public health, but also in terms of odours or insects.

CRM: But do we not have resource conditions that stipulate what is acceptable there?

WWM: Yeah, but everything in the plant is covered by consents. Isn't that the same for water?

CRM: In terms of compliance, water quality, yes that's what I'm saying

ESM: I think the difference though is that everyone knows about when the water quality goes wrong. But if you have a big discharge somewhere, and if you don't tell anyone, then probably nobody even knows about it. That's completely different...

WWM: Yeah, I agree with that. But if we've got midges people know about it. If we've got odours then people know about it.

CRM: Yeah, now you have said to me previously that those are the two things that Mark is concerned about. If we get complaints in that area, it illustrates two things. One that you wouldn't be complying with your consents, and second that the company spent \$500 million and didn't solve that problem, so that's a potential issue of stakeholder confidence.

WNM: Mhm, what about where you have a short term unconsented wet weather overflow into a public area? There's huge consequences with that compared to the same thing occurring in a stream where no one notices. Like when we had a wet weather overflow into Kelly Tarltons...

ESM: That's flooding...

CRM: Well that's the other thing that dawned on me, [WWM], in relation to wastewater. My perception is that we could never have a wastewater failure as catastrophic as the most catastrophic water failure. Is that fair comment?

WWM: Yeah I could agree with that...

WPM1: Come on swap salaries then...

ESM: I think the thing with wastewater is that the private individual always has service. You can always flush the toilet and it goes somewhere.

WNM: There's never service interruption...

CRM: [WWM], if I asked you to formulate the risks you're most worried about at the plant using this framework, what would be the impact of say a centrifuge failure? Cos basically I'm saying that it would have to result in a resource consent non-compliance or a business inefficiency. Or what it is it? Is it...

WWM: Or stakeholder confidence. The last three columns would probably cover it. But say if we lost all centrifuge capacity then we would have a bloody mess.

CRM: Yeah, but I'm thinking that in order for you to drive development at the plant, you have to quantify your needs in these terms for the Capex application...

WPM1: Well, if you took a hard view of the world then you could argue that Capex's at Mangere might not exactly be flying through the system, unless you justify it on the basis that you've spent so much money to get our discharges clean that we have to maintain that level of expenditure. Cos there is another option. We could just stick it all out into the tide, so what. But I guess the customers have said they want a higher level of environmental quality. You have to judge it.

RS: So it's a stakeholder expectation, isn't it, as well as our own internal expectations.

WPM1: They say they want a clean harbour where they can gather shell fish. Well then, [WWM] needs centrifuges to achieve that.

WWM: Well if they failed you'd have to bypass the plant and feed sewerage straight into the harbour...

CRM: So it would be a non-compliance.

WWM: It would be a non-compliance, there would be health issues...

RS: It's an unacceptable level of performance.

CRM: Yes but the rationale here is that the health risk is managed by the resource consent conditions. So, in other words, the reason we have a resource consent condition is to protect public health.

WNM: Or do resource consents provide for environmental health? I mean do you have a health regulator and an environmental regulator, [WWM]?

RS: You have to notify Auckland Regional Council for breaches, but do you also have to notify the Health Protection Officers at the local councils?

WWM: Yeah there is a requirement under certain conditions.

RS: So there is a public health expectation...

CRM: That's after you've breached a consent is it?

WWM: Yeah.

CRM: So the presumption is that provided you never breach a consent, there's never any health issue?

WWM: Well...

CRM: ... My fundamental question [WWM] is are we fairly representing your primary drivers if we tie everything in terms of the wastewater risks back to compliance with resource conditions, stakeholder confidence, and impact on business efficiency?

WWM: Well, yeah probably, but that's just looking at this. We won't know until we use it. And anyway, we've got a risk management system at the moment. What are you trying to do with this? Make it more transparent in terms of the differential between the different risks?

CRM: What I'm trying to do is to provide the parameters that you guys would consider to be your bottom line, the things that are fundamentally important to your business operations. So when I talk to Mark the things that he always talks about are grading issues, compliance with the bulk water agreement, and health and safety. Resource consents are the other things that are consistently coming up as being fundamentally important to the business. So what I'm trying to do is provide you with parameters that you're more familiar with on a day-to-day basis, that are more tangible to your business operations.

WWM: Well you know I know what my ten greatest risks are and this doesn't change them. Is this being used to try and make them more transparent?

RS: What we're really trying to establish is some measures that are realistic in terms of your particular areas of the business, in terms of the consequences of failure. At the moment the questions in the framework are pretty broad, subjective questions. We're just trying to remove some subjectivity and provide a scale that makes it easier for the asset management people and for yourselves in terms of scoring, ranking, and prioritizing risks.

CRM: The aim of changing the framework is to make sure that those ten risks that you know should be better represented by using a framework which is better suited to your context. So in other words the ultimate aim here of changing the framework is to get better quality data, cos at the moment the quality of our data is hindering the asset management group. It's all about getting an improvement in the quality of the

data, at a very basic level it's about getting a better representation of what you know. And what I'm asking you today is will you get a better representation using these parameters here? Am I providing you with the right vehicle to reflect what you guys know in your experience, in your knowledge?

WPM1: Well I think yes you probably would. But not if you continue with the theory that a catastrophic resource consent breach is the same level of catastrophic as mass boiled water notices, mass dirty water, mass no water, or public health sickness. I think you have to slide these up so that they all sit at some higher point. And you might find this unpalatable, but I think the same goes for financial impacts. If we lose money, so what? The company would survive.

CRM: The whole idea is to ensure that when you evaluate your risks you don't get some bollocks result that you know isn't real. So when [WWM] says I know what my top risks are, and he uses this framework to evaluate them, and he looks at the result, he says "well, in general, yeah they're about in the right order". Whereas at the moment I'm concerned...

WPM1: Well then this needs to be skewed towards water quality and water delivery.

CRM: I don't think there's any doubt about whether that will occur.

WPM1: So as you say a resource consent failure can't hit a 100.

CRM: A resource consent failure, yeah ok that might be possible.

ESM: Yeah I think you need to push all the overflow stuff down the chart a bit, cos overflows are, well wet weather ones anyway, are part of normal operating business.

WWM: Ok, but what's the overflow from the plant?

ESM: That's a different category.

WWM: So where does that fit?

ESM: That's non-compliance, so in the middle there, you should have ongoing minor non-compliance.

WPM1: Is the overflow from the plant a minor non-compliance or a large non-compliance?

WWM: So we're saying there can't be anything at the treatment plant that could be catastrophic? Because the only thing that can be catastrophic is not meeting water quality standards, which is fine, so that means that anything that happens at the plant is only very serious?

CRM: Yes that's correct.

WWM: So a major failure at the plant, and we pour raw sewerage into the Manukau harbour for two months...

ESM: That could be catastrophic couldn't it?

CRM: Well it depends...

ESM: It would be a prominent breach...

WPM1: I think you could have a catastrophic failure at the plant, but its rating would be say 90% of a catastrophic failure in the water network. So you could say yes it's a catastrophic failure, but it get's weighted to come out less when you apply your numbers.

WWM: I think this is what you were saying before. It probably means that the overflow stuff needs to be scaled down so that you can be up there, but its not...

CRM: Not quite as high as a serious water quality interruption...

## Dialogue No. 5 – Engaging with stakeholders: Project Managers

**Source:** Transcript 13 June 2007 (CRM:02)

**Participants:** Corporate Risk Manager (CRM)  
Manager Project Services (MPS)  
Various Project Managers (PM1 – 4)

---

CRM: ...and through that process I realised how absolutely challenging it is to use the risk register. We basically had a workshop and went through and tried to identify risks and then I sat down after the workshop with [PM on Project Hobson] to have a go at scoring them, but I realised very quickly that it is extremely difficult to actually capture the impact of project management failures using the current risk scoring method. Now I guess I've been here for long enough now to understand that risk management's really been driven by the Operations group for ten years, and over time the development of the whole risk management function has been more and more pushed towards suiting their needs, so the business has kind of ended up with a one size fits all solution...

PM1: Yeah I suppose we try to use it for the drivers for the Capex, but not really for the management of risks during the project...

CRM: Right, so I'm starting to introduce the idea of having different systems in different areas of the company so that each system suits the needs of that business unit, but all feed into a common risk classification system. So hopefully in a month or two from now we can have a project risk register that's actually more in line with project management work, but that requires me to understand your work environment, and what kind of dictates the way that you work. And I guess one of the obvious things that differentiates your group from a project management consultancy is that you are actually part of a wider business so you have to worry about the fact that your downstream customer is also your business partner. So ultimately what I'm trying to get from you today is an understanding of your work flow process, and then I'm going to go away and have a look at things and come back to you with a proposed register.

MPS: By the downstream customer you mean Operations or...

CRM: Yeah I'm assuming that your customer is the Operations group. Now I know we have the Project Delivery Manual, but at this stage I don't have a very good understanding of your work process, of where you get your work from, at which stage you take over, where you take it to, and at which stage you hand it over to

Operations, and what's involved in that. So I guess the first thing is what's the interface with Planning, does all your work come through the Planning team?

Several: Nah, no...

MPS: Largely...

PM2: Well it should do but Ops tend to come to us through the back door. They push a lot of condition refurbishment work, pump replacements, and the like to us.

CRM: Ok so Op's sneak in primarily rehabilitation work?

PM2: They really should be going through Planning because Planning have a strategic view of our systems, and say if we want to upgrade something Planning might say that in two years time we won't actually need that asset anymore cos we've got a bigger one being built next door. So strictly speaking Ops should go through Planning, and then Planning get the CAPEX, and then we carry out the design and implementation. That's the way it should work in theory but it doesn't...

PM1: Well also the limit for Opex is \$20,000. But if they have something over that then it's deemed to be a project so we pick it up. And the Networks guys are short of staff so they're happy to wash their hands of anything over 20 grand cos they haven't got the resources.

CRM: Ah so primarily small stuff you mean?

PM1: Anything over 20 grand. It might be something worth 50 or a 100 grand but its really just an ongoing maintenance task, maybe some big paint job or something, which you could argue is that really appropriate for us to do. And from our point of view a lot of those would be more efficiently dealt with in the Operations group because we have got relatively high overheads. You know we have to follow the Project Delivery Manual, and the approvals process, and then there are the interfaces, managing operational constraints, and the handover back to Operations. For us those are quite big overheads, but if it was done in Operations a lot of that stuff can be short circuited.

CRM: Right ok, so that's a bit grey, the stuff that comes through that way.

PM1: Well there's also been quite a bit of movement in the interface with Planning over the years. Before the last restructuring it was quite clear that it was up to Planning to get the Capex and then they would hand it over to Projects. But since the last restructuring the line has moved. I mean, certainly we need to have input to the Capex cos sometimes the Planning guys got the dollars wrong, but I think some of us feel that the line has moved too far since the last restructuring. We end building all of the business case as well which really should be coming out of Planning. A lot of the stuff we're getting from them is just too high level...

CRM: So it's a level of detail issue?



PM1: Yeah, we're getting to the stage where all we get from Planning is the justification, so such and such asset needs to be replaced. But that's not the project, I mean what are we actually supposed to build. So a lot of the investigation and development of the solution is coming over to the projects side, which is actually causing some difficulties with us because we need a lot more lead time for that and we don't have the investigation budget for it.

PM2: Yeah, really the investigation should be thorough enough to give us all the information that's necessary for a Capex request. But what we're getting is back of fag packet jobs, straight lines on a drawing, "there you are guys, that's what you need to do, get on with it".

CRM: Right ok, and yet the guys in planning presumably don't want to work at that level of detail?

PM1: Probably, I don't think they currently have the resources to get down to that level of detail either.

CRM: Ok... so in theory when you get a business case that has been approved to go ahead as a project, that approval should occur via a Capex?

Several: Yes

CRM: Ideally that would be delivered by planning but sometimes you're finding that you have deliver it.

PM1: Well no. We are responsible for putting up the Capex's. I sign all the Capex's, but the business case is supposed to come out of planning.

CRM: Ok.

PM1: But they say that something needs to be replaced, and a big chunk of the business case is actually evaluating the options and deciding which one to proceed with. The trouble is that I think Planning are having less and less involvement in that.

CRM: Ok. So when do you normally get involved in a project?

PM1: Well normally we have to get involved with the preliminary design and the investigation, because otherwise we can't do the Capex properly.

CRM: Alright and following that what's the next stage?

PM2: Detailed design and development. What we should have at Capex stage is a partly investigated solution to go forward with, and we take that solution and develop the design, take it to tender and then construction.

PM1: Our first lead may be getting the consultant on board, especially for staged Capexs where stage one would be for the detailed design. Then we would have to put a brief together, get the consultants' proposals in and evaluate those, and then get a preferred proposal approved before we could proceed.

CRM: Do you use consultants for the investigation work at the start as well?

PM1: We may use consultants for those early investigations. Some we do, some we don't...

PM2: Usually goes out to competitive bids. That's the way we work. We have to go out and get some competition on any work that we want done.

CRM: Depends on the scale of the project, presumably?

PM1: Well [the CEO] frowns on having the same consultant doing the investigation, and then cos they know what is happening having them do the detailed design and supervision. So often we have to change horses at each stage.

CRM: Right, ok.

PM1: So virtually everything is by competitive tender, although [PM3] managed to put a package together for the bulk supply points. He put up a delivery strategy which varied from the norm, and proposed using the same consultant throughout.

CRM: So a preferred provider.

PM3: Well the first task in the PDM is to prepare your project delivery strategy but often the strategy is unwritten. Its usually just agreed amongst the [Project Control Group] that we're gonna go the traditional route or otherwise. Vary rarely do we actually write down a formal project delivery strategy, but the project [PM1] referred to is one exception.

CRM: Is that because of time constraints or because you don't normally see value in preparing the strategy?

PM3: Well normally we just follow the PDM process which really has the standard delivery strategy already mapped out. But on that project we wrote a decent paper on what we intended to do cos it deviated from our norm and we wanted buy-in at a higher level to make sure we were going to get the necessary backup at the key stages.

CRM: Ok. So after the detailed design and development you put it out for competitive tender, we choose a contractor from that, and then I understand that you quite often engage someone to do the site supervision. Is that right, we don't do that ourselves?

PM3: Sometimes.

PM4: Either we get somebody to do the site supervision or we use Don's team.

MPS: Because we don't have enough staff to supervise all the contracts, and some of them are of a specialist nature, so we need additional expertise.

PM1: And also the size of the programme. Very big projects can tie up all your resources.

PM3: And also the geographic distribution. We've got jobs in the Hunua's and we've got jobs in the Waitakeres, and it doesn't make sense to send one PM to two jobs that are geographically split like that...

MPS: You'd spend half your day on the road...

PM3: So all those factors need to be taken into account.

CRM: Right. And after the asset's been built on site, there's a commissioning phase?

PM1: Well we have a process with Operations in terms of planning on the way through for taking plant or pipelines out of service or for doing cut-ins. We have monthly sessions...

PM3: Yeah we have a shutdown co-ordination programme which lists all the potential or real shutdowns that we need to plan for, so we can make sure they are all co-ordinated so we don't have clashes.

MPS: And it's also subject to other events outside your control, like the weather or yesterday's Hunua No 1 break.

PM1: Yeah, it was not even a year ago that most of this stuff was just in the heads of guys like [PM3], the experts. But as the headroom in water treatment and networks has gotten less and less, the co-ordination of shutdowns has become more critical. And I think Operations have become more risk averse, or perhaps more aware of their risks, so now we have a sheet that we fill out for a shutdown which we give to Operations so they can approve it or not, although how many of those do we get back signed?

PM3: I've not seen any come back to be honest. [laughter]

CRM: Right ok.

PM4: Yeah there's a register for shutdowns. I think it's established procedure at the moment...

PM1: Yeah that's right. It's a significant management of risk.

PM3: Last year's the first time we did it. Myself and [PG] put together a little bar chart and we kept track of what we had planned, when it was going to happen and how each one would interact, so we could look ahead and see that say if that one doesn't happen next week then we're in trouble for the following three weeks because we've got these other shutdowns, so that one's gonna bump a month if it doesn't happen next week. So it just enabled us to collectively manage the project risks. It was a huge step forward, and then this year [PM2] done a lot of hard work on this programme so we've taken it a lot further and got fine piece of work now rather than just a coarse management tool.

PM4: Yes, and now it covers networks and treatment.

CRM: Ok right. So after commissioning presumably we have things like as-built drawings and manuals and things. Do the project team deal with that?

PM1: Yup. We typically have to put together a commissioning plan...

PM3: And all the testing as well, the pre-commissioning work. Often it's the testing phase, testing, disinfection, cleaning, whatever you want to call it, it's quite a lengthy and highly detailed planning process to get to the stage where we can actually switch some pumps on actually run something. And in fact the difference between a pumping station and a new piece of pipeline is quite radical. The commissioning is quite different. With a pipeline, one day it's not online, we do a shutdown, we tie it in, and then it's online, it's commissioned whether you like it or not.

MPS: Open a valve and she's on.

PM1: And the other thing is that we have to get the assets in the assets register. Cos the assets might be in service but unless they are in the assets register and Mosaic they can't log the maintenance records.

CRM: And is your team involved in putting those assets into Mosaic?

PM1: We prepare everything that needs to be uploaded, but Operations does the input.

CRM: Ok. And does your team deal with the LNOs at all?

PM4: Yeah we do...

PM3: Yup.

MPS: Not directly. Only when we're actually doing some of their work for them.

PM3: Well we do where they are a key stakeholder in the project. In that case in stead of a PCG meeting, Project Control Group, we have a PLG meeting, Project Liaison Group, and that will involve the LNO or some other key stakeholder who's not normally part of the overall management team.

CRM: Who else can be involved in those?

PM3: Could also be the Council rather than the LNO, also Iwi...

PM4: Auckland Regional Council, the parks...

PM3: Roading and parks departments, Transit New Zealand.

CRM: Ok that's good. Now in terms of using the existing risk register, what are some of the problems that you've had with it? What's your opinion of the register?

PM3: Well we've all been involved in developing the current version. Six or seven years ago we were doing very little in terms of formally recording project risk, it was pretty much an ad hoc process. That was one of the things that I highlighted when I joined the company, there was no real vehicle at any point in the project for having a register of risks to use as a management tool. But we were constrained to what was in the existing corporate framework, and that's made it quite messy, having to calculate the risk class. It would probably be more appropriate to say it in words rather than trying to put numbers on it. But that's just my opinion.

CRM: Like moderate, low, high?

PM3: Exactly, because that's all we use it for anyway. At the moment Class 1 or Class 2 risks just get broad brush mitigating measures, things we tend to do anyway as part of the general management process. We really only focus specifically on the Class 3 and 4...

CRM: So let me ask you, do you normally do different things for those risks, or do you...?

PM3: Generally they are quite specific. It could be that you need a formal service agreement with Transit New Zealand and there'll be a 28 point list of conditions associated with those. For other's Communications might be a major risk so we will prepare a communications plan.

CRM: So it is driving some the stuff you do then.

MPS: The other thing is when you go through and score a risk, whether you do it numerically or in words, and it ends up being a Class 5 you think "shit, can't have that", so you end up massaging it to make the five disappear. Which I guess is really where it tends to fall down.

PM4: Yeah, it's the same when you have the feeling that certain risk needs to be a Class 4 cos it is important so you make the numbers right that it gets a Class 4. So the numbers don't really help a lot. The classification is useful, but tweaking the numbers to get the right classification doesn't really help.

CRM: Right ok I understand, using the framework to match the result to your intuitive feeling for the risk.

PM3: Yeah. I think high, medium, low is a much more palatable way for us to go forward. It would save a lot of time as well, cos a lot of effort is required to populate all the fields in the register and you end up with a quite a chunky document.

CRM: And I'm aware that there's a vastly different scale of risks between a \$50,000 job and something like Project Hobson, and the register doesn't really capture that.

PM1: Yeah, on a \$50,000 dollar job chances are most of the risks are stuff all anyway, compared to the scale of the Hobson risks.

PM4: Should we do a risk register for a \$50,000 job?

Several: No

PM4: Where's the cut-off point actually?

PM1: Well it's not formal is it, that's the trouble.

PM2: We sort of agreed that it was Board level didn't we, so \$2 million?

Several: mhm, yeah

PM1: Well we do the formal register for the Board reports, but we've got a generic register for smaller jobs.

PM3: That's right. Basically the generic register has got eight or nine categories with a handful of generic risks in each one. They may or may not apply to every project but generally they're a pretty good scoping document to start from.

CRM: Ok alright. Now I guess the fundamental thing is this, how does Watercare actually feel the impact of project risks, when things actually go wrong? I mean, I'm under the impression that often if there's a delay in a project it doesn't have that big an impact on the business.

PM3: It's just a cost at the end of the day.

PM2: Generally not. It depends if there is a deadline to get a service. If there's key dates obviously we try to meet those, but otherwise a slight delay mightn't have, be an issue for us.

PM3: No, it'd just be an extension and an associated cost.

PM2: Depending on what causes it.

PM1: Most projects come through Planning and are driven by growth, so we have to get the upgrade or the new facility done in time to meet the growth in demand.

CRM: So in a way Planning would have a huge impact on whether time delays are critical or not, if they've got their act together and they give you a large window.

PM1: Well Planning often don't consider the delivery side of projects closely enough, particularly timing over the summer months. But we're hoping to work much more closely with them on the AMP this year.

CRM: Ok, and Operation's? I mean, I'm presuming that if you deliver a rubbish asset, or if the asset wasn't well designed or something, then ultimately Operations feel the pinch?

PM1: Well these guys hear about it for years. [laughter]

PM2: We generally bring them in at the start of the project, get their buy in, and get them to review our preliminary and detailed designs, and they certainly get to see the tender documents before they go out so they can have a good understanding of the project and what they are going to get.

PM3: The Project Control Group always has at least one person from the Operations team, so they are party to the monthly meetings, and they get to review the preliminary and detailed designs and the tender docs as [PM2] says, and all of the key milestones that are involved, so we don't end up delivering something they didn't want.

PM1: We get Ops signature on all the drawings so we've got that buy in at the end of the day, and we've got project handover process that's detailed in the PDM. The problem though is that the information doesn't always get disseminated fully amongst the people within the Operations group.

- CRM: Ok. I mean that's how you're mitigating the risk, but what I'm trying to gauge here is how the business would ultimately feel the impact of something going wrong on a project.
- PM3: A big driver for buy-in really is that in order to deliver a project we've generally got to do a shutdown of some description, and Ops will not shut the network down unless they're confident that we are ready to implement the final part of the project solution. Cos these shut downs generally happen in a one day, eight hour, six hour window, and the confidence level's got to be there that the project team's ready to do what they say their gonna do. So you've got that buy-in level as well.
- PM2: Well projects are unique and the risks associated with projects are totally different than our standard risks. By the very nature of them some of the risks can be very high. I mean look at Hobson. What's the worst thing that could happen? The bloody tunnel could collapse I guess.
- MPS: And that's the nature of tunnelling projects, that's the kind of thing that can happen.
- PM2: But it's big on the project scale of things. It doesn't really affect Watercare. From that perspective the outcome is just that we don't get the asset online until later.
- CRM: Yeah, and that's what I'm trying to get to terms with...
- PM1: Well it depends on what impact you're trying to show. You can have a massive impact on the project, but perhaps not so big on the whole business. Say if you're putting a pipeline through, sue you don't want to go killing someone, but if there's a bit of a collapse and you lose a bit of pipe and it costs a \$100,000, what's that in the big scheme of things?
- CRM: And that's the thing. The more I think about it, it seems to me that you guys don't have a critical impact on the business, at least in the short term.
- PM1: Well you can have a very critical impact when you're doing a cut in...
- PM2: And if there are critical dates...
- PM3: Yes, we could impact significantly at shutdown time. If there's a fuck up during a major shutdown then you will see impact.
- CRM: So that's the difficult part for me. How do you guys have a big impact on the business?
- PM1: Well there isn't much when you're working offline. I guess the biggest risk is the chance of killing or injuring someone, cos of the nature of construction work.
- CRM: Ok, let me set something out to you. I think that the risks that exist in Projects, a bit like Planning, are more long term. So there's probably stuff all in your area where something could happen tomorrow and have a big impact on the business. But thinking say five or ten years down the track, the things that could put us in a really bad position are really under the control of Planning and Projects. I think there is a

need to consider the significance of risks that may not be realised for five or ten years.

PM3: I'd agree with that. But there's probably one exception and that's the reputational situation. We can cock up quite badly out there and cause a serious impact the business, say by causing traffic grid lock day after day after day...

CRM: Or killing people

PM3: Yeah or flooding property

PM4: Or when we are making critical changes to the system. The work we're doing at the moment changing the chlorine dosing lines at the Huia Treatment Plant is a case in point, cos chlorination of water is a Class 5 risk at the Huia plant.

PM1: Yeah, there's another issue about how we look at these things. Now I know we have to manage these things as project risks, when we do shutdowns, cut-ins, commissioning and that sort of stuff, but I wonder whether those risks should really be in the Operations register. I haven't thought this through in detail, but if we cock up a cut-in and no-one gets any water, sure Projects have done it, but really it's an operational risk. There are a lot of things about the operation of the plants and network that we can't control in that regard.



## Dialogue No. 6 – Engaging with stakeholders: Asset Managers

**Source:** Transcript 2 July 2007 (CRM:05)

**Participants:** Corporate Risk Manager (CRM)  
General Manager Planning (GMP)  
General Manager Operations (GMO)  
Manager Project Services (MPS)  
Project Manager (PM1)

---

CRM: ...the project managers seem to think in terms of time, cost, and quality if as the key parameters, but the score risk is pretty well tied to the operational context, so trying to quantify project risk using the current scoring framework just doesn't seem to work very well. The obvious one is that it's very difficult to capture project delays in the current register so there's a bit of a fudge factor that's used to do that.

GMP: You can't turn those into dollars? I mean, do they translate into liquidated damages?

CRM: Well...

GMP: I guess there's two components to it isn't there? One is a delay, and the other is the overhead costs that we have to pay if we extend the contract.

PM1: Well we have got the dollars there because we have to do the calculations that go into the contracts. But often they're not that large are they?

GMP: No.

MPS: No, on an enterprise-wide basis they won't be very large.

GMP: In which case the risk is low.

CRM: So I've been trying to get a handle on how to evaluate project risks. Ultimately I'd like to strip the project risks out of the register and hold those completely separately to operational risks. Project risks will include Planning and Projects, because the more I look at it, the two processes can't be separated. So I was looking at the corporate objectives that drive projects, and those seem to be either compliance, the need to satisfy future growth demands, security of supply, customer demands, or service level improvements...

GMP: Replacement too...

MPS: That's level of service or maintaining level of service rather than enhancing level of service.

GMO: Ah, the other reason is to save money, where it's just a straight economic decision.

GMP: Not many of those, but occasionally there is one.

CRM: Right ok. So they then set the objectives of the project in terms of what the project's got to deliver, what I'm calling the functional objective, and the time scale in which you've got to achieve it, and then when you've worked out what resources you need to achieve those, how much the project is going to cost. Which explains why we've got cost, time, and quality as the three key parameters. But then how the organisation actually feels the impacts of risks at the project level comes back to our ability to comply with regulations, to satisfy growth, or perhaps we don't deliver the necessary service level improvements, or we don't get the efficiencies we wanted, or we end up carrying more risk than we would like. And all of those are inherited by the operational group in the end aren't they?

MPS: Well other than Greenfield growth. Cos at the end of the day if you don't get the water supply or wastewater to a Greenfield site then it's not an operational risk like if you hadn't put in extra source capacity and therefore don't meet your drought security standards.

CRM: Ok. So the first thing I've been trying to work out is how Watercare actually feels the impact of risks in your area, and other than perhaps cost over runs I think the impact really lies with the Operations group who inherit the results of those risks So whether or not a project represents a major risk to the organisation, and it seems obvious to say this, but it seems to depend on the criticality of the project to the corporate objectives.

GMP: Yep.

CRM: And so I guess what I'm struggling with is how do we measure how critical a project is to Watercare? Because I've got to tie the assessment of risk at a project level back to impact at a corporate level, and the impact at a corporate level depends on how critical that project is.

MPS: Depends why you're doing the project...

CRM: So if you have a three month delay on a project, does that actually matter at the corporate level? It may not unless that project...

PM1: Well often the time scale is so long it doesn't matter.

GMP: Yeah, generally for projects there's something wrong if the time is tight.

MPS: Correct.

GMP: Generally it wouldn't matter if it was two or three years late. And it shouldn't matter because we have a 20 year AMP so we should know about projects well in advance of when we need them. Something has gone wrong at the planning level if we need to suddenly have a project delivered tomorrow.

MPS: Correct.

CRM: Well looking at that from my point of view, having such a long lead time for projects is a way to mitigate the risks.

GMP: Yep.

CRM: But I'm assuming that you looked at it from [GM Finance] point of view he'd like to say why can't we just leave it for another three years and deliver it just in time.

MPS: Yep.

GMP: Agreed.

GMO: And that's where projects like Hunua No 4 get quite tricky...

CRM: Because it's a security of supply risk?

GMO: Yeah. You can say well we've lived with it for 10 years in this mode, another couple of years won't make any difference.

GMP: Generally speaking, we could stop all of our work for a year or two and...

GMO: And nothing would happen...

CRM: You wouldn't notice...

GMP: Well, theoretically nothing would happen.

PM1: But then gradually it starts to get worse. You get to the stage, like in some cases now where you can't do upgrades because you've left them for too long and there's not enough headroom left.

GMP: Yeah. If you've got it right then you should have time. You should have an even work flow, it shouldn't be all peaky. There's a reputational risk to the organisation if we end up with large peaky work flow and we then struggle for resources. But then a lot of projects are to mitigate identified operational risks, and really it doesn't matter so long as their implemented before those risks come to fruition. Which is where it's difficult from a timing perspective, how do you build the Hunua No 4 just before the Hunua No 3 fails?

CRM: But from a planning point of view is that not challenging? Because at the moment you just have Class 1 to Class 5 risks. You don't have a line that says this one's acceptable but this one's unacceptable, or where you'd be able to say this risk is going to become unacceptable in five years time.

MPS: Well I guess the priority would tend to be all of the operational risks that are already in that Class 4 or bordering on Class 5. Those and the growth projects are probably the one's that have got our attention.

GMP: Theoretically, if you do nothing over time then the risks of all of our assets should be climbing and what we're trying to do is take the one's that are getting close to Class 4 and 5 and put them back down into the Class 2 and 3. But over time they will continue to grow back up there.

MPS: Yeah that's right...

CRM: And are we limited in the scope of projects that we can tackle by the amount of funding that we can request. I mean if we went to the Board and said "look we've actually found we've got a hundred bloody projects rather than ten projects that we want to tackle in the next five years, give us a hell of a lot more funding...?"

GMP: We've got an example of that at the moment where we've pushed the AMP and it's caused a reaction. For the last few years we've had a fixed price regime with our customers and as a consequence we're going to get this spike in pricing. So I guess that's an example of it having not been done well. If our planning had been better six years ago then we wouldn't have that problem now.

GMO: At the moment the Board are basically saying that, putting aside the cost of inflation around construction, the AMP is now a ceiling and we're not allowed to go above it.

CRM: Sorry can you explain that one again?

GMP: Well if we go back six year's, the planning wasn't working well in the Water part of the business and as a consequence we didn't identify that right quantity of work coming up. As a consequence of that we entered into a three year process of fixing prices, CPI adjusted. Now about three years ago there was a change in how planning was done and we recognised that there was a historic problem in the Water planning and we've been addressing that over the last three years. So last year's AMP had a significant increase in spend forecast on the Water side of the business, and I'm not talking about 20 years out, but within the first five years. As a consequence of that we've had to signal that price's will have to go up a long way, and that's causing the organisation all sorts of problems.

CRM: Ok.

GMO: A related issue there is that I think our tolerance for wastewater risk is probably slightly greater than for water risk, though I'm not sure to what extent our register reflects that at the moment. In other words if it's your last ten dollars and you can either fix a water risk or a wastewater risk, I suggest you fix the water risk because the environmental issues are not seen to be as critical as the public health issues around water supply.

CRM: Ok.

GMP: For example in the UK they are capital constrained. They enter into an arrangement with the regulator that restricts how much money they can spend. So Yorkshire Water were telling me that they have this five year programme of work, thirteen thousand projects on their books and they basically ranked them so they could draw a line depending on how much money they got allocated.

CRM: A cut off.

GMP: Yeah, so they are ranking their projects to tell how far they can afford to go. We don't really do that. I guess we do it on a kind of informal basis. We're trying to manage our ability to deliver projects, so we try to smooth the amount of money that we've got to spend over time.

CRM: Right ok I see

MPS: We've not really had that hard internal debate over how you do that ranking. I guess if you had some capital constraint in terms of total funding available that would force that debate in terms of where would it be best to spend those sorts of dollars.

GMP: Yeah.

CRM: So what you're saying is that we're not capially constrained but we can't throw out sudden peaks in the work flow.

GMP: Well we are at the moment.

PM1: Well project delivery is constrained by what we've put in the AMP isn't it?

GMP: Yeah, the constraint is really the impact on pricing.

CRM: Ok. So from my point of view the issue really comes back to the point that was just made about how do you go about determining the criticality of projects if you've got multiple projects and they're satisfying different objectives.

MPS: Correct...

CRM: So growth, compliance, and things like that. Presumably some of them are a higher priority than others. I guess one of the examples would be the works done to satisfy the change in drinking water legislation two years ago. Presumably that became a critical project if you like for the organisation.

GMP: Yeah. And that's an example of a project that did have an end date that was quite important for the organisation.

CRM: So I need to come up with a framework that captures that criticality so that I can convert the work at the project level to an impact at the corporate level. Now 90% of the time I don't think we're going to have projects that are turning up as enterprise critical, but at the same time, I don't want to give the impression that project work is not important, that there's no need for a project risk register because it won't add any value.

PM1: So there's two different things, is that what you're saying? There's a project risk register to try and minimise and mitigate risk within the actual project itself...

CRM: Correct, and then I need to convert those project scale risks to an impact at the enterprise level.

PM1: Yeah but they're different aren't they? Those risks we have within the project are sort of within the project loop, which is different from whether the whole project...

CRM: Is significant, that's right.

PM1: Yeah, that's actually the one that rolls up, not the individual project risks.

CRM: That's right.

GMO: So with that risk-based approach can you actually use that to calibrate our prioritisation process with the AMP? I mean, if you take the list of projects in the AMP, effectively, for better or worse, we must have prioritised those and said well these are more important than those...

CRM: Hmm that's right you're doing it intuitively.

GMO: Could you actually go back and look at those risk scores and see if you could make some sense out of them as to...?

CRM: The idea really is for me to capture what we're doing intuitively and the discussion today has been quite good for me in that sense. The other challenge is that at the moment the projects team are required to measure their risks using a scale that's designed for Watercare as a whole organisation. So that measures financial losses in the order of millions or tens of millions, which means that losses at a project level often turn up as Class 1 or 2 risks. Ideally, to give these guys some resolution in their project risk register we should have a smaller scale for measuring consequences for projects, or perhaps for different projects. So a million dollar project would have a different risk register scale to the one that's being used for Project Hobson. But what I don't know is whether or not we have some kind of financial scale system already within Watercare that I could try and tie into. I know that there is the Board approval level of \$2 million, which is an obvious cut off, so perhaps projects over \$2 million use a different risk scoring system to projects under \$2 million. Are there any other kind of steps in financial scales that we use?

MPS: Why do you want to look at that Jason?

CRM: Well, going through a risk management process, going through all the risks and scoring them is meant to give some guidance on where you should be focussing you're efforts. But if the manager on a small project is measuring everything on the scale that we use for Watercare everything turns out to be very small, and there's no differentiation. It also implies that those are insignificant risks to the organisation.

MPS: But say if the worst exposure on a contract was \$50,000. If our enterprise risk scheme says that's insignificant then why re-classify it.

CRM: Well if we don't then we're kind of saying that it doesn't really matter if we're not efficient at the project level, cos on a project by project basis those sorts of losses it won't impact the organisation.

PM1: But that's why you're saying we need a separate project register, because it should actually be used as a management tool by the project managers.

CRM: By the project manager's, yeah.

MPS: Yeah, but don't you risk ending up with a different register at every level of the business, because \$50,000 for trade waste might be important or \$50,000 for stores might be important, whereas enterprise wide it's not really that important. Aren't you in danger of getting people to focus their efforts in the wrong areas.

CRM: There is a balance required.

GMP: The reality is that losing \$50,000 on a job is important, and we should be tracking that in terms of the performance of project managers. Maybe it needs to be in their PADR, or some other sort of measure completely outside the risk process.

CRM: Yeah it's possible. From my point of view, it's an issue that hasn't been resolved. It's a matter of coming up with the cut-off, or perhaps having some sort of variable scale. I guess these are the issues I'm trying to tie down through this process.





## Dialogue No. 7 – Engaging with stakeholders: General Managers

**Source:** Transcript 15 October 2007 (CRM:15)

**Participants:** Corporate Risk Manager (CRM)  
Various General Managers of Watercare (M1 – 4)

---

CRM: ... in the foreseeable future I'm aiming to have three frameworks in place, and by frameworks I mean the way that we evaluate risks. At the bottom will be what I'm calling the business operations framework. In essence this is concerned with looking at how we might fail to meet our performance requirements given the capabilities that the business already has. What I've tried to do is to determine what our performance requirements are from the Statement of Corporate Intent [SCI], which has about 18 different strategic performance targets. I have clustered them into five fundamental categories that I think align with what the business fundamentally exists for. So, on that basis, I'm saying that the five fundamental performance requirements for the business are "to provide water", "to treat water", "to meet statutory obligations", "to operate in the most economical way that it can", and "to keep our stakeholder's confidence".

M3: Ah, you're talking about "delivering water" and "treating water". Why isn't there an equivalent one for wastewater?

CRM: Well when something goes wrong in wastewater operations the performance failure ultimately boils down to non-compliance with a resource consent. So the performance objectives for wastewater are really to comply with resource consents or statutory obligations.

M3: I think there's an inconsistency there because, in the same way, there is an obligation to treat and deliver water according to the public health and drinking water standards.

M1: Yes, where you have "delivering water" and "treating water" I assumed that "treating water" was treating wastewater. In other words "water" was just being used in a more holistic sense.

M2: Yes, I did too.

CRM: Well, bearing in mind that we're talking here about the content of the Statement of Corporate intent, what would you say are the performance requirements for the wastewater operations? When we treat that wastewater, we are treating it to what standard?

- M1: To a safe standard for the environment.
- M2: Yes, treating it in accordance with our obligations under the resource consents, in the same way that we're treating water in accordance with our obligations under our contracts and the drinking water standards.
- M3: Perhaps one thing to consider is that if there are anomalies in the SCI, why not use this process to try to align them?
- M1: Well the SCI is a bit of a camel to be honest. Over the years the SCI has morphed into sort of a dumping ground for ideas. Seems like every time some one comes up with something it goes in there, which is frustrating to manage. And because of that it's not a particularly good SCI. It's got a lot of extra stuff in it.
- M4: Yes, it has a lot of tactical and prescriptive requirements.
- M3: If you come back to the basics of why Watercare's here, it's to deliver water to an agreed standard, and it's to collect wastewater. They're the two bottom lines of why the business is here. You can't really have one without the other, so I think wastewater does need to be incorporated somewhere.
- M1: Yes, that's right. I've always looked at the company has having two customers. One is the LNOs [Local Network Operators], or rather the people of the region through the LNOs, and the other is the environment. Actually, why can't we just restate our Sustainability policies? Don't they capture the organisation's objectives?
- M4: Rather than the SCI, look at the LGA [Local Government Act] and the objectives which are specified there. They are basically maintaining the integrity of our assets, delivering services, and operating at least cost.
- M3: The other one that I keep coming back to is when we were doing the Three Water's Strategy we spent bloody ages thinking about what were the drivers for this. And it came down to Growth, Regulation, and Levels of Service. And if we're going to get consistency then we somehow need to bring all these things together and agree the word's that we're going to use in future documents. I would suggest that in the next month, you should try to get together all of the documents that have got something to contribute here and see if you can rationalise them. So the LGA, the SCI, Three Waters, the Sustainability policies. Let's see if we can get a group that we can all work with in our documents and get consistency.
- M4: I agree, it's a bit of a mish-mash at the moment.
- CRM: I need to reiterate that risk is about performance. So you need to have a defined objective for what you are trying to achieve, and then you can measure risk against the achievement of that objective. So in the end we need to boil it down to a couple of fundamental objectives which apply across the enterprise. If we want to measure risk across the enterprise they need to be the things that fundamentally drive it.

Now, by the sounds of it, I have mistakenly assumed that the SCI was that document. It certainly reads that way.

M3: Hmm, to be quite honest, I thought the SCI would be one of the primary drivers of everything we did so its quite interesting to hear the explanations today.

M2: Well you can't ignore it.

M4: But it doesn't drive the business fundamentally.

M2: Well the driver, fundamentally, is the legislation. That's fundamental. I look upon the SCI as a bit of a translation, even if not a very good one, but a translation of that into some practical measures. So what does the legislative objective mean? It means, among other things, that we will operate to a one in two hundred year drought security standard. It's trying to give some measurable words to the general statement that we will provide water and wastewater services to the region. So they're not disconnected. But the legislation is a more permanent and wider base.

CRM: And that's an important point. We have to be able to provide quantifiable measures for the performance that we're measuring in each area. We have to keep that in mind. It's very challenging, and that's why it takes so much time to get it right.

M1: Well, there's no single answer. What we're talking about is which boxes do we break up the Watercare world into. There's a whole lot of different ways of slicing and dicing the business, and many different answers. We need to agree a common group that we're all comfortable with that traverses the full breadth and depth of the business.

M4: I actually think that where you have here [reading from CRM's notes] "the five fundamental business objectives representing an aggregation of the strategic performance requirements", it should be driven by the strategic plan not the contract with the shareholders. Because the contract with the shareholders does not drive the business. What drives the business is the Board's enterprise strategy. That's actually at the core of everything, rather than the SCI. The SCI represents a contract between the business and the shareholders, but its not a strategy. Now, we have got a strategic plan that's needs renewal. It's been five years since we went through a strategic planning process. And I've been trying to put together a framework to take us forward with that, and it seems to me that it should be risk based. Because we've got the Three Waters strategy that needs to feed into it, and we've got the SCI which is the shareholder's component, and the AMP [Asset Management Plan] and Funding Plan, and then the director's views about where they want to take the enterprise. It seems to me that the unifying piece is really the risk framework. Maybe we need to think about how the risk framework is going to drive the strategic plan.

M3: Well that's interesting because we're facing new challenges with the Local Government Act and the four well-beings, and I think in the future we're going to driven more and more by that through outside parties as well as the Board of

directors. And I'm trying to take that into account in the areas under my responsibility. In the Three Waters strategy we had an options evaluation process based on the four well-beings. And Jason and I have been talking about trying to categorise strategic risks under the four well-beings, to try and get some of these linkages. I think that fits in beautifully with an integrated strategic plan.

M1: Actually, the four-well beings really drove the development of our Sustainability policies and objectives. I took all the information there was on the company, cut it all out into bits of paper on a big table, and effectively parked everything up in a home. That's how we ended up with those six policy areas, which, in a way, are our interpretation of what the four well-beings are for Watercare. They're just another way of talking about the four well-beings.

M3: The important point is what you start with as your driver for all these things, and I think there's a debate that needs to be had about whether it is the four well-beings or something else. What would be helpful, and I think this will be critical for the strategic plan, is to start with what drives the strategic plan. Which presumably is the SCI and the LGA, they both have an influence on it. What else? What drives that? And how does it fit with your HR strategy, your information management strategy, the sustainability policies? Does that diagram exist within the firm at the moment?

M4: Yes, it's published in the Annual Report and the AMP, what we call the business planning cycle. But I don't think it's quite right. It indicates that the SCI sits above everything else as the driver of the strategic plan, but that's not right. Essentially we have mistaken what the core starting point is, and it's the enterprise strategic plan that should be the starting point. We just need to refresh it, to get it right, and there's an opportunity here to have the enterprise strategy dovetail with the enterprise risk framework. I think there is an opportunity through this process to identify some category four risks at the enterprise level that haven't really been considered, like resourcing and to have a discussion about those. And if the strategic planning process is working well then it should identify those areas where we're not necessarily planning well, because we tend to plan for assets and roles and capital. So what's missing?

M3: Yes, I think it's also a factor of the stage of maturity of the business. You've got a lot of things under control which has been the prior focus. No you need to look beyond that. Energy is another one in my view.

M4: I agree. While we have been historically good in certain areas, there are others where, while we are developing them, there is a sense that there's no vision. The core systems replacement is a good example. We're about to completely replace our main business system, but we have no top down view of what we want that system to deliver. Do we want a real time end-to-end system, like Vector, or do we want to look like Manukau Water and just have an asset management system without much integration. There's no vision. I mean the organisation really does have

a planning gap on the people side and the technology side. It's just a bit piecemeal. So there is an opportunity to shift the resource development strategy into a better state of clarity.

CRM: One point that I do need to get across is that the risk management framework is not a chapter within the business, it is something that applies across the entire business in terms of how we measure our needs. So when we say this is the strategic plan, it's not like we can say well risk fits there, it actually sits across and at all depths.

M1: I agree.

M4: So do I. Where risk needs to develop is actually at the strategic level, that's where the missing link is. And I see that you will help us to develop the strategic plan through an appropriate risk framework.

CRM: The practical challenge is to marry the strategic stuff, the high level drivers, to what guys on the coal face talk about on a day to day basis. There's a real challenge to provide practical content in a framework that is designed to meet strategic needs. So that's where there's a lot of time needed to develop this.



## Dialogue No. 8 – Reflecting on the dialogue with the General Managers

**Source:** Transcript 19 October 2007 (CRM:16)

**Participants:** Corporate Risk Manager (CRM)  
Richard Donnelly (RD) [*Researcher*]

---

CRM: So I've had a look at the business objectives today because that really should be the starting point for measuring performance. There seem to be really two principal sources for the corporate objectives. First, the Local Government Act [LGA] defines what we're required to do and that is to provide water and wastewater services, we must "manage the business efficiently with a view to maintaining prices for water and wastewater services at the minimum levels consistent with effective conduct of that business and the maintenance of the long term integrity of the assets". We're also required to deliver an AMP and a Funding Plan. Second is the Company Constitution, which says the principal objective of the company is "to operate as a successful business which provides water and wastewater services that are economically viable, environmentally sound, socially responsible, and responsive to customer needs". Those seem to be the driving documents. Basically the LGA says the company exists to provide water and wastewater services through an efficient business, and the Constitution says we have to do it in a way that is economically viable, environmentally sound, socially responsible and responsive to customer needs. So there are four categories there in terms of how we will do things. Now those have somehow been translated into a series of objectives under six categories, the Sustainability policies, which to me is not great in terms of alignment with the primary ones. To me there are four obvious categories there for how you operate your business, but we've used six and I'm not sure why. And what's really interesting is that there are two primary documents that detail our company objectives, one is the Annual Report, and one is the Statement of Corporate Intent, and the body of objectives in both of those documents are different.

RD: Yes, I noticed that too. Has anyone explained that one to you?

CRM: The way that someone put it to me is that the company objectives listed in the Annual Report are basically a bottom up construction of business objectives. In other words the business units have said "these are the things we are conscious of", or "we could provide these measurements of how we've improved things", and they have built up the objectives on that basis. The same person also suggested to me that the Statement of Corporate Intent was a top down construction. Now I don't necessarily agree with that, I think realistically the Statement of Corporate Intent is a

wish list from our customers and shareholders. The point is though, that neither of those bodies of objectives have been constructed as a deconstruction of the fundamental corporate objectives, which then begs the question of whether either of those documents is actually comprehensive, let alone the fact that you can't reconcile them. The same person also suggested to me that the Annual Report is a backwards looking document, where the Statement of Corporate Intent is a forwards looking document. So the Statement of Corporate Intent supposedly defines what we will do as opposed to the Annual Report which kind of summarises what we did and how successful we were. I guess you could look at things that way, but to me the two documents should be reconcilable. I think there's a major risk to the organisation that you publish a document and say "this is how we've performed against our objectives", and then someone picks it up the legislative document that's meant to define the company's objectives and says "hang on, these don't match".

RD: Yes, well you would expect the Annual Report to report on the objectives in the SCI, and that both should reflect the legislative objectives.

CRM: Exactly, so I think there's an issue of concern there. I think what's actually going on is that, practically, the content of the company's Annual Report has been driven by a desire to keep the report leading edge so that it continues to win awards, and that process has not been linked to updating the SCI. So in the end it's not really driven by the company's objectives, which is a bit arse about face. And that's a problem for me, but I think I can put together a case for a project to rationalise those objectives. In that regard, if I just recognise that [the SCI] is a customer wish list and [the Annual Report] is in essence a marketing thing, then maybe I don't really need to pay too much attention to them in the interim. They sure as hell don't reconcile so why should I be too worried about it? But the LGA and the Company Constitution are fixed in stone as far as I can see, and it seems that the legislation defines what we do while the Constitution indicates, to some degree, how we will do it. So we will provide water and wastewater services, that's what we will do; we will manage our business efficiently, that's how we will do things. Other obvious ones: maintenance of long term integrity of assets, I think that's what we will do; we will prepare and supply an AMP and Funding Plan, that's a bit of both, or at least supplying an AMP is what you will do, but it's also a process that you go through in order to deliver water and wastewater services. And then economically viable, environmentally sound, these are all kind of how we will do things.

RD: But they're all reasonably vague terms. What "environmentally sound" actually means has not been specified.

CRM: Yes, there is definitely flexibility in the sense that the standards we are required to perform to could change, how we measure things could change, and even the scope of activities that fits under the umbrella of water and wastewater services



could change. And that makes it difficult for me. How do you measure the performance of the business? I mean there are lots of performance requirements here. There are lots of commitments, things the company has agreed to do, or is required to do, and there's too many for me to build into a practical framework. And some of them I need to separate out. For example, "effective conduct of the business", that's an all-encompassing statement; and the other one, "maintenance of the long term integrity of assets", I think it was probably meant to be physical assets, but it could be interpreted in a broader fashion, as in knowledge. So how do you define and measure those commitments? I need to work out how the company measures whether it is economically viable, environmentally sound, socially responsible, and so on, and then those measures presumably go into the risk framework. And unfortunately I want to get there quickly now. I need to take these different commitments and condense them, cluster them into a manageable number of consequence measures, and then have indicative measures of what would constitute different magnitudes of failures. But it's proving difficult and in all honesty I think I'm more confused than a few weeks ago before that meeting [with the Risk Management Steering Committee]. It's not clear in my mind. I think there's a differentiation between what we do and how we do things, so there are two sources of risk for us, risk in terms of trying to deliver what we're supposed to, and then risks in terms of trying to operate as we're supposed to. But I don't quite understand what I need to have to capture those.

RD: By 'capture' you mean...?

CRM: In terms of quantifying those risks. The thing that stands out immediately is that if you want to quantify them then you need some measure of performance, so some measure of service delivery standards or requirements, and some measure of business process or management requirements.

RD: Right, so the question is how do you group those commitments, and then how do you represent each of those categories? I mean, the measures you choose will effectively represent those categories, so represent risk, and what you're trying to avoid is a situation where the measures are ambiguous, which would cause confusion about how to evaluate risks.

CRM: Yes, and that's something that's confusing. I'm not sure how many frameworks I need. You remember how with the project risk register I defined two contexts, the corporate context and then the project context, so the risks get calculated for both contexts, and they're linked by understanding the corporate objectives for the project. You can apply the same theory to business units in the sense that you can have an enterprise that has four different business units, and if you could define the corporate objectives for the business unit, then you could set up a different framework.

RD: Yes, I guess you could.

CRM: And I think that would be a lot more efficient in the sense that the framework would be much more specific to the business unit so it would be able to provide the manager with more useful risk information. I mean, if there was just corporate level framework and all the managers had to work off that, then they would lose a lot of resolution. And also that can send a really bad message because often things which are important at say the project level or the business unit level are not so significant at the enterprise level, and that lack of resolution can send the message that those risks are not important. So I think there's a potential advantage in defining specific frameworks for different contexts, so maybe for Finance, Asset Management, Operations, Business Services. If you could define what the specific objectives were for each of those units and show how they contribute to the overall corporate objectives then you could presumably provide them with risk frameworks and risk data that are more pertinent to what they do. But I suspect that trying to be comprehensive in defining what the corporate objectives are for a business unit would be very difficult particularly where business units work across the organisation. And then there's a potential issue about what happens if the organisation's objectives change?

RD: Yes, if the risk framework is defined by deconstructing the objectives, then the framework would need to be redefined if the objectives changed.

CRM: Exactly, it's an interesting point. I mean Watercare's objectives seem to be fixed, but what about organisations that need to be very flexible and react very quickly to develop and pursue new objectives. If you've gone through setting up all these kinds of frameworks, then you may have to change those frameworks when you decide to change your tactics. So what quality or what depth of information should you develop the risk frameworks to if the organisation is one that needs to be flexible and change direction regularly?

RD: Not only that, but what about performance standards? Watercare's objectives may be fixed at the highest level, but as you said before, the standards to which the company has to perform could change. That would imply a need to at least change the performance measures in the frameworks.

CRM: Right, so that's another issue of concern, fixed objectives at the very top, but potentially fluctuating objectives at lower levels. It would seem to beg the question, how can risk management be effective? I mean, is risk management potentially always lagging behind the organisation? To keep the frameworks up to date you have to invest huge amounts of money, and then by the time you've done it they're potentially out of date, so you have to ask are we getting the returns on it? It does suggest that to be practical your risk framework needs to be tied to very very high level objectives. You really need to have something that's not going to change. You've got to have the fundamentals of the business, what things are we always going to be doing?

- RD: But that then begs the counter question of how, if the framework is detached from the specific objectives, then doesn't that violate the fundamental definition of risk as uncertain effects on performance? Although, having said that, maybe this issue explains why, in the commercial sector, risk is typically defined in terms of a single performance measure, so financial return or economic capital. I mean, in the financial sector risk is categorised according to the nature of the causal event, so there's market risk, credit risk, and a huge number of operational risk categories, and all of those risks are measured by just one indicator. So there, in a sense, you have a single high level objective and performance measure which doesn't really change. Whereas what we're talking about here is categorising risk according to the nature of the effects on different dimensions of performance. Maybe that's an important difference with public sector organisations, the mission objectives are just as important as the financial objective.
- CRM: Well I think it's an important theoretical question because I've realised that one of the main reasons we use risk data is so that people can make an assessment of business need considering the corporate objectives. So, at least in theory, you can look across the entire enterprise and say "this risk data gives me a basis to decide what to do". It's an objective method for determining the priorities based on what the organisation wants to achieve, not what I want to achieve. So, to me, you have to have a consistent and objective method for calculating risk, otherwise you can't compare across the organisation.
- RD: Well yes, where you have incommensurate objectives competing for the same pot of money, like we do here, that really emphasises the importance of clearly defining the corporate objectives.
- CRM: Doesn't it, screamingly, because what you're saying is that the corporate objectives will in effect define the relative importance of different things. And that's something I'm very conscious of at Watercare. If I develop the risk framework using the objectives currently under the Statement of Corporate Intent, then I think it would mean that we would stop spending money on wastewater, which would be completely unpalatable to [the wastewater side of the business]. I mean, in essence I'm saying that all risks will be assessed on the significance to the achievement of corporate objectives, but the objective most directly relevant to wastewater is that there will be "no successful prosecutions under the Resource Management Act". The problem is that it says "prosecutions", it doesn't say "compliance with resource consents". And, what's more our resource consents at the moment are under the consideration of the [Auckland Region Council], and practically I think there's fairly good agreement that they're almost indefinitely on hold.
- RD: Yeah, it potentially won't be resolved in even 20 or 30 years.
- CRM: That's right, so the probability of us getting successfully prosecuted for wastewater overflows is probably very small. In which case, why spend money on it? I realised,

holy shit, if we asked [the Wastewater Treatment Plant manager] to do a realistic assessment of the chance that we will get prosecuted, its very small, and his risks will come out Class 2, and he'll never get any budget. And so, what's going to happen is he's never going to endorse the risk management function, because he'll recognise straight away that he's got no business case, his business unit becomes insignificant. As a cultural thing that's really dangerous. It's a classic situation of unintended consequences. What the objective should be is compliance with statutory obligations. That's a very different objective which would place wastewater on an even or at least a comparable footing with water.

## Dialogue No. 9 – Justifying the risk framework

**Source:** Transcript 23 November 2007 (CRM:18)

**Participants:** Corporate Risk Manager (CRM)  
Richard Donnelly (RD) [*Researcher*]

---

CRM: I'm writing this document which I've described as establishing the risk context. In essence it describes how I've arrived at the structure or the methodology that we will use to quantify risk. I guess there are probably three principal reasons why I produced it. In order of priority for me, from a personal perspective, one is to rationalise my thoughts. I find that by writing things down with the objective that other people will have to read them I am forced to try to structure my thinking, and I guess you could say that I am forced to analyse my intuitive conclusions, or forced to justify my intuitive feelings. In some ways that's making explicit what's implicit in my head. So that's the first reason. The second reason is that I figure that by doing that, by better structuring and more efficiently communicating my thoughts so that they are less muddled, the [General Managers] should hopefully be able to more easily read it and understand. So it's meant to provide them with a reference document which can act as a basis for them to discuss things with me. And then the third thing is as a knowledge retention device. I mean, I am sure that [all the previous Risk Managers] went through this process of defining how are we going to analyse and quantify risks here at Watercare, but none of them left a document which outlines the rationale behind it. And now I'm having to rework or recover that old ground. My thoughts are that when I leave this organisation, ideally, someone else shouldn't have to go through that. Within the first week they should be able to pick up the document, read it and be able to short cut, presumably, what took me three to four months to understand.

RD: I know there has been some sensitivity in this organisation to you wanting to change the risk framework. Do you think the document will help to overcome that?

CRM: Well it's interesting actually. When I started writing it up I thought it was very black and white, but I've realised through the writing that are definitely points where what I've done could be disputed, where my conclusions could be questioned, they're not absolute statements. So, in some ways, the document is more about justifying my result than it is about documenting how I arrived at the result. It's really a story of how I view the company, or how I've interpreted the company's objectives. To give you a practical example, I analysed the customer contracts for the supply of bulk water and concluded they boil down to two primary issues, water quality and water delivery. But when I think back, I knew a long time ago that I wanted water delivery

to be one of the ways that we should measure consequences. I guess I could have gone another way, I could have sliced the contracts from the [CFO's] perspective and concluded that they're primarily concerned with how we charge the customers, or maybe I could have come up with ten other things. But the reason I came up with water delivery was because I was conscious of the fact that one of the ongoing problems here is that we need to be able to predict future risks, and the solution that I arrived at within about six weeks of being at Watercare was that service delivery performance can be modelled. Always in the back of my mind was this thought that we need something that we can model, we need to have a measure of consequence that we can model, and since we already hydraulically model the water networks I wanted that to be a measure of consequence. And so my summary of the customer contracts, in that sense, is probably biased. But that's not to say it's wrong. Bias is a bit negative. The other way you could put it, which is more constructive, is that I am aware of a wider picture in terms of taking in to account how we can satisfy our risk data needs in the way that I interpret the company's objectives. And ultimately that could be the major selling point. The way that I'm slicing and dicing the objectives, or rather the way that I'm analysing the organisation will enable us to use computational models to do some of the work. For some reason, that's what will probably seal the deal with the General Managers. Quite why that should seal the deal, really, when you think about it, I don't know. But I'm confident that it's like my ace card. It's a major advantage over the current risk framework, and if I can get them to hold their objections until I get that across then I think they'll agree it's a good idea.

RD: That's interesting, you're explicitly linking the risk framework to the analytical capabilities in the company. The benefit being that by doing so you get analytical data to support risk assessments, rather than subjective judgements as with the current framework.

CRM: Yes. But it does beg the question, just because we can use computer models, does it mean that I've cut it the right way? I mean Watercare has an AMP process which prioritises the company's projects. At the moment that prioritisation is all pretty much intuitive, but effectively that process is already doing what I'm trying to capture here. I'm just trying to put a more formal assessment process on it, to say that this objective is actually more important to the organisation than this one. It's an awkward question for the organisation to deal with, and even do you need to, I think, is the bigger question. I mean, that's what's rolling around in the back of their eyes. There's a perception on their part that I might be trying to put mathematics onto something that they do already. So to be honest it's a fair question. Do you need to do this when there is already a process in place to do it? When does it become too much?

RD: I don't know, but certainly linking to those modelling capabilities adds more support to your framework.

CRM: I don't know either. But I think I can argue against the way the current framework is set up. I think it comes back, and I've talked to you about this before, to the fact that I've been trying to marry the operational focus with the strategic focus. So I started with the objectives and tried to break them down, but I always thought that at some point I love to find that they marry well with what the guys on the shop floor think about. So I've talked to them and worked out what's important to them, the point really being that if we're doing computational modelling then its probably computational modelling of what they think about, isn't it? So my awareness of what we model at the moment has probably influenced the way I've dissected the organisation's objectives. As opposed to the current framework which I think was really an arbitrary categorisation of the strategic objectives, and they've gone like this [makes dumping motion with hand]. Now that's fine, that's a representation of sorts, but there's clearly a disconnect between the way the framework represents the business and the way they think about and analyse it at the coal face.

RD: That's an interesting representation you used, the analogy that they've gone like that [makes dumping motion with hand], that the framework they created only represents the organisation from the top-down perspective. Whereas, what you're trying to do is to marry the bottom with the top. If you look at those two approaches, then the existing top-down framework doesn't take into account the capabilities or the resources that are needed to convert what people at the coal face are thinking about into the terms of that framework. So, and you've commented on this before, each assessment becomes quite subjective, depending on how the person doing the assessment interprets the framework. Now what you're saying is that you've created your framework with an eye toward the analytical capabilities that already exist, so that those capabilities can be used to make a good translation between what the guys on the shop floor are thinking about and what gets represented at the top.

CRM: Yeah. That's that whole concept of having objective criteria. I wanted our risk frameworks to be constructed in a language that the guys at the coal face use on a day to day basis, parameters that are familiar to them. So what I've done is tried to find out how they are thinking on the shop floor, and tried to see if I can deconstruct the strategic objectives in a way that is consistent with that. And then as a result, perhaps not surprisingly, what we find is that the company already has a lot of the analysis in those terms already. So we've already got the RCM models, the hydraulic network models, the compliance monitoring system, the public health risk management plans. So the irony is that the selling point seems to be this computational modelling. I guess you could say it stands as evidence of the effectiveness of the translation of between the coal face and the Board room. So if

the risk framework is a communication conduit between the coal face and corporate then, in essence, the fact the categories in my framework reflect well the business systems that we already have in place is evidence of how efficient the communication should be. As you were saying, it's difficult for [the Networks manager] to convert his knowledge into the framework because the translation is so subjective. Whereas, with the way that I've constructed the framework, there will hopefully be less translation required.

RD: Or perhaps a more accurate translation.

CRM: Yeah, a more accurate translation.



## Dialogue No. 10 – Feedback on strategic risks

**Source:** Transcript 21 September 2007 (CRM:14)

**Participants:** Corporate Risk Manager (CRM)  
Richard Donnelly (RD) [*Researcher*]

---

RD: ... what's happened is that they [the managers] have identified a range of general issues which are or may potentially cause some uncertainty. But they were pretty vague in their description of those issues. Often they have just written one word or a vague statement, so like "terrorism" or "tightening up of standards". But they haven't given any specifics, so, for example, there's no indication of what standards they're thinking about. Even if they have given a specific like that, they then don't go on to link it to what impact they imagine the change might have on the business.

CRM: Right, ok, so no basis if you like.

RD: Yes, they haven't really justified why they wrote something down. So there's potentially a question around how many of these are just general concerns that are always there versus things that someone actually knows there's a specific reason why they need to be concerned about it now.

CRM: Right, so an elevated reason, or trigger to be concerned about something.

RD: Yeah. Also, your instructions were that they should think about issues on a long term, 10 to 50 year horizon, but mostly they did not give any specific indication of the time horizon they were thinking about for each issue. The exception there was [the Chief Information Officer]. He said that given the nature of the IT business, it was not meaningful to consider a time frame longer than 10 years. So, perhaps with the exception of climate change and the long-term energy situation, I think most of the issues they identified are realistically located in the short to medium term, mostly less than 10 years.

CRM: Really?

RD: Yes, but I think that was probably to be expected. It's an issue commonly encountered in scenario planning that it is difficult for people to comprehend just how much things can change in even twenty years, let alone fifty years. It's difficult for people to step outside the box of their day-to-day activities and think about what issues might arise 20 years from now. They're necessarily concerned with issues that are probably already on the horizon so that's what they tend to write down.

CRM: Right, so looking at this concept of strategic risk, have you had any ideas on how to go about taking this information and making a strategic risk evaluation? Is it as simple as saying these are some of the issues that have the potential to impact on

our business or on the environment in which we are practising and what are the implications of that for us? I suspect that's probably how its done in a lot of cases, but to me that just smacks of a brainstorming workshop, which I hate.

RD: Well, that's what you've got with the data there. They've effectively just brainstormed a whole bunch of issues.

CRM: But I mean in terms of risk identification I've been advocating a more structured method. So taking objectives and working down, but I'm not sure that you can do that at a strategic level. I guess you could do a SWOT analysis of your current strengths and weaknesses versus the environment. So we could some predictions about what the future environment might look like and then define our current state in terms of strengths and weaknesses, and then try to evaluate that if that was the environment then what would we do, what would we want to do, what would we want to be doing now in anticipation of that environment? So you could say this is our current state, we know that's definite, and these are the various states that could exist in the future, those are highly uncertain. There are various ways the future can branch, so which ones are more important, and how can we maintain flexibility to those futures?

RD: You're talking about scenario planning.

CRM: Yes, do you know much about how to do scenario analysis?

RD: Only what I've read. Something that is emphasised as important though, is the need to have a single focussing question. They always address a specific question, so perhaps there is a big investment and the question would be what are the potential future conditions, and what does the investment look like under those conditions.

CRM: Right.

RD: The other thing, in terms of evaluating the future environment, even qualitatively, is that the list of issues we've got here are very interconnected.

CRM: Yeah, that's something I'm just picking up going through them. Working at this level it seems that we're looking at the interrelation of lots of factors, a network of influence.

RD: Right, so in terms of framing scenarios I think you would need to group these issues according to some of their commonalities, and then construct scenarios around each of the groups. So some of the groupings that seem to stand out are around procurement, so trends in the costs and scarcity of materials, consumables, and infrastructure services, leading to heightened supply chain risk; around human resources, so concerns about the aging workforce and skilled labour markets, increasing labour costs, and problems securing the right people. Another one is concerns about the increasing complexity of technology, and also resource planning, so concerns about the future development of Auckland, and about

maintaining the quality of raw water supplies. A big one related to that is of course climate change. There were quite a number of climate change related issues there. I think these kinds of groupings imply certain sets of focal questions that could be addressed with scenario analyses relatively independently of each other.

CRM: Yeah, at the moment I'm not sure. I need to sit down and think about what I'm trying to do. I certainly want to have some assessment of strategic issues or risks. I think it comes back to this thing that if I want to identify strategic risks then I have to ask "what are we trying to achieve at a strategic level?" and then "what could stop us achieving that?" The problem at the moment is that as far as I'm aware there's not really any definition of what we're trying to achieve at a strategic level. I mean in private enterprise strategy is about where do we need to take the business, in terms of expansion or new markets or something like that. But there doesn't seem to be any equivalent here, it seems like Watercare exists for the purposes of existing in a way. Perhaps it's the nature of the business, being in essence a local government authority. Maybe that's just it, the business can't develop like that, it just needs to maintain itself.

RD: Well, fundamentally the mission of the company is fixed by the legislation. It doesn't change. The only thing that changes are the performance objectives in relation to that mission, so how much, or what standard of service.

CRM: So then when you talk about strategy then, what is strategy? Is it strategy for the way that you deliver your business? We could ask "how is Watercare going to deliver water in the future? What's the strategy for delivering water in the future?"

RD: Well, you're not in a business where you can choose your market, you've got your market, and you've got an infrastructure system that you have to work with. It seems to me that the strategy then is around how could that change in the future. Now, you've also already got the planning department which is concerned with evolution of that system. That's what they're there for, to plan how the system is going to evolve. Since that is the case, perhaps strategy should be looking at the fact that in order for them to do that job they have to specify a whole lot of assumptions and parameters that define the future environment, so assumptions about the demand profile will be, assumptions about the standards of service. Strategy could be altering those parameters and those assumptions to see what effect they have on the planned development of the system. So, for example, one that was brought up several times in that list is a concern with the effectiveness of regional landuse planning and uncertainty about the movement of businesses off shore. Now, at the moment, the demand profile which is used for planning purposes is based on the population estimates produced by the Auckland Regional Council, and the demand estimates from the individual [Local Network Operators]. But if there was a lack of control in the urban environment in terms of land use planning and unforeseen or unaccounted for migration of businesses overseas then that would fundamentally

change the demand profile. So perhaps there are variables out there that Watercare should be worried out in terms of the effect on planning parameters, which could be addressed through strategic scenario planning.

CRM: But isn't that different. To me management is about how you do something and strategy is about what you are trying to achieve. The way someone put it to me was that say there is a group of people lost in the bush and one guy climbs up a tree and says "that's where we're going" and then the other guys, who are the managers, they basically take the machetes and cut their way through the bush. So the guy who stands up the tree is setting the direction, which is the strategy, and then the other guys are managing how to achieve that. So to me strategy is all about direction as opposed to what you were saying...

RD: Well, there's almost two levels of direction here. One is the guy up the tree is telling them that's where we are going, but he's also got to map out the route and tell them what's in the way and what the obstacles are so they know what they're having to deal with, cos if they're on the ground they can't see what's in front of them. So it seems there's two levels of strategy there, one is the objective, and one is the things we need to take into account in order to get there. That would be my analogy for Watercare's situation. The fundamental objectives are pretty much set in stone. The strategy is about working out how to deliver them.

CRM: Yes that's true. Right so there's the goal if you like and then in between is the path [drawing on paper] and then it's a question of how. What you're saying is the goal is pretty well established.

RD: I think so. It's set in legislation so...

CRM: Yeah, so unless the legislation changes the goal is set. So in other words Watercare has to deliver these services in the future, and so strategy is about identifying the obstacles that management are going to have to deal with. So a strategy for Watercare, a strategy document might just simply be saying that if we're going to continue to do business under our current remit these are the future obstacles. So those are the strategic risks?

RD: I would say so. They are sources of uncertainty, or sources of uncertain effects on the business, if we use Ward's terms.

CRM: [Drawing] So for Watercare, and maybe it's just this company, may strategic risks are those big ticket things that will affect the company's ability to deliver its purpose. Alright ok, it was never going to be simple unfortunately. It really does show how you have to think about these things though. It's only when you stop and think that the problems become apparent. You know, I thought it would be relatively straightforward to ask people what were the strategic risks facing this organisation, and they say well its climate change etcetera. So I need to think about how we define and identify strategic risks.

RD: Well it's interesting that even though they identified a number of things which in a generic sense seem like strategic issues, virtually none of them wrote down what the impact on the business might be.

CRM: No, but that's become apparent though. In this business a risk is an asset, but the actual concept of risk is about the consequences and the likelihood of those consequences occurring. Certainly one of the logical things is to ask well what could cause it to happen, but the most important thing is to define the consequences and to ensure that when you assess the likelihood you're really assessing how likely those consequences are, not just the likelihood of the event. But that's okay, its a start. It's really a matter of having triggers to prompt people's thoughts.



## Dialogue No. 11 – The purpose of “risk data”

**Source:** Transcript 21 September 2007 (CRM:14)

**Participants:** Corporate Risk Manager (CRM)  
Richard Donnelly (RD) [*Researcher*]

---

CRM: ... I'm developing this [risk assessment] framework now, and I'll have to sit down with staff and say 'this is your new framework, and I want you to apply it'. But when I ask them to apply it, in order to be credible, I have to be 100 percent sure about what I'm doing. When they ask questions I have to be crystal clear, I have to have pre-thought about everything. And that's one way in which these discussions with you are actually proving fruitful for me. But it does make me wonder about how often people in risk management roles actually get the chance to sit down and think about what it is that a risk management system is supposed to be delivering. Take this for example. Question 14 from my survey was 'how accessible is the corporate risk register for you to use?' Now I looked at the response data, [reading from survey report] "the distribution of responses to question 14 indicates that 2 out of every 3 respondents believe that accessing the corporate risk register is problematic." Ok, so staff can't access the corporate risk register, but [reading from notes] "an important element of integrating any process into business-as-usual is ensuring that people have the information that they need." Question 14 was intended to assess whether or not the accessibility of the corporate risk register was restricting staff access to information. Currently the corporate risk register is maintained by me, but it's freely available to all staff on the intranet. But this, [reading again] "the other notable fact is that sixty one percent of respondents currently don't appear to use the risk register. This could imply one of two things: firstly risk data is not being used broadly within the business or alternatively that some staff are sourcing data through other staff members." Well, the second part there, that's not so bad, but this first one, "risk data is not being used broadly within the business". So there's not a need for risk data? That scared the willies out of me. That one sentence horrifies me cos it says 'what's the purpose of my role?' So why are we doing risk management? Well, when I realised that, when the penny dropped, I said [reading from survey report] 'there's a need to clearly understand where risk data is being used or could be used to improve performance within the business. This understanding is necessary to ensure that the format and scope of risk data satisfy business needs.' So the question is when do we use risk data? At bottom, what is the purpose of risk data? The only current needs that I know about are the formulation of the Asset Management Plan [AMP] and capital expenditure [Capex] applications, so I started thinking about these. The AMP and Capex applications are both concerned with

change, so the risk data is being used to justify change, it's evidence in some way. The risk data is evidence for change. And I was trying to put this into my little framework for how I see the business, so [drawing on paper], over here we have business Operations, and over here we have the Project delivery group. The way I have interpreted the function of these two sides of the business is that Operations is about working with current capabilities, while Projects is about delivering new needed capabilities. It seems then that risk data is used to justify change, to justify the projects that the company undertakes. So I was thinking about that, and then I started thinking at a more basic level, what are you doing when you assess risk, what are you looking for? You want to know the higher risks, the large risks. Why? Well, if we're talking downside risk, then these are threats to performance and you're trying to identify them, and to measure them. So that then is the risk data, the information that feeds into a decision on what to do about the risks, how to treat them. Hmm, you know, this is just going to bring me back to the [Risk Management] Standard isn't it? That's the basic risk management process isn't it?

RD: Ah, yes. If you continue that diagram then you'll end up drawing out the steps of the basic problem solving process, the basic risk management process from the Standard.

CRM: Hmm, yeah. So what are we doing? We're generating risk data as a basis for decision making? So it's a knowledge base isn't it? And the decision making is concerning what? It's decisions about your performance, I think... [drawing] reliable performance, reliable achievement, more reliable achievement.

RD: That seems consistent. The risk management theory talks about better decision making. Risk data supports decision making.

CRM: Hmm, but it seems to me that there's an easy case for risk data here [pointing to Business Projects], risk data is used to justify business projects. But there's not really a demand or a use for risk data here [pointing to Business Operations]. Risk data isn't used in exercising your current capabilities. But if you really wanted to get value, if you really want to so show that risk management was valuable, then you would find a case for using risk analysis here [pointing to Business Ops].

RD: Sorry, can you explain that? I'm not following that risk data is not used in Operations.

CRM: Well, I mean, Operations is about your current capabilities, the things you do over and over again, the regular repetitive tasks. Projects are about changing your capabilities, they are less frequent, perhaps more important, but I'm not sure about that. And, in order to work out what capability you need for the future, you need to understand what your current status is. At least in Watercare, projects, I mean the need for a new capability is driven by a realisation that your current capability is going to be short at some time in the future, or its not going to be adequate, or



something like that. So somehow you must be analysing what you're currently doing to recognise that need.

RD: Yes, that's logical, how else do you identify what you need to do?

CRM: Well, isn't that what risk management is doing? The risk assessment is asking 'what is this organisation trying to achieve with its current operations, what could go wrong?' So, for example, we've got to deliver water, what could go wrong to stop us delivering water? So we do those risk assessments and then if the risk is too large, if it is unacceptable, then that triggers a project. And it's the risk data that is used to justify that the risk is unacceptable. But the point is that we are not using risk data here [Business Operations]. It seems like Operations is where we collect all the underlying data, and we use that data to make an assessment, but that seems to be the only place we're employing it, to justify change, to justify a project. We don't actually seem to be using that data in Operations. The data gets extracted into Planning and then used to justify projects. What I'm saying is, can all that data we generate in Operations actually be used in Operations as well?

RD: Hmm, but what about Reliability Centred Maintenance? They have used risk as variable there. That's an example of risk data being used in Operations.

CRM: Yes, you're right. Let me think about that. What are they doing? In essence they're taking risk data and they're using that risk data to influence their decision making about maintenance. So they've identified risks that could result in kit breaking down, they've got information on how the kit could break down and the relative importance of that, and then they decide how to manage that kit so that it doesn't break down. Is that a form of treatment, or management, as opposed to change? So in other words, is that decision making as opposed to change? Now I don't know that the two things are different, do you know what I mean?

RD: Ah, well, I'm not sure that they are different. Perhaps we can look at this a slightly different way. What we're talking about here could be described as an issue of defining the type of decision we're looking at. So, at least intuitively, it seems easy to identify something called a decision when we're talking about spending twenty million dollars on a new asset. There's a long planning process, there's a formal paper trail, and there's definitely a point at which someone puts a signature on a piece of paper and says 'Ok I agree with the business case, here's the authorisation to go ahead and do this.' So it seems easy to classify that as a decision. But it's much harder, or it becomes increasingly more difficult to identify the actual point of a decision as you go down in the organisation into decisions that happen more and more frequently. Eventually you get down to everyday decisions that people are making about how to do their jobs and deal with the particular problems they're facing at the time. And I think what you're talking about is that the use of risk analysis, risk data, seems to be limited to the really big, clearly defined decisions in

the organisation. And the question is, what about all the others? Why can't risk management play a role there too?

CRM: Actually that's very interesting cos one of the things that worried me in the survey, and I'm not sure if this is exactly the same thing, but one of the questions in the survey was [reading from survey report] "do you know the risks that may adversely affect the performance of your business unit?" And there was a distribution of responses which indicated that [reading from survey report] "greater than fifty percent of respondents have no formal knowledge or awareness of the risks that could impact the performance of their business unit." I wrote [reading from survey report] "this is a concern because it may indicate that the staff follow business processes because they're supposed to rather than because they understand what the business processes are designed to achieve in terms of risk control."

RD: Yes I think that's what we're talking about. If you take the idea that risk is a measure of a problem, then if more than fifty percent of respondents don't know the risks for their area of the business then they don't know the problems that they should be dealing with.

CRM: Or for which they're delivering a solution. That's right, 'I don't know the risks but I know that I do certain things a certain way'. Ah, so that's the key I think. I've forgotten about how risk management acts a control function. Just intuitively, this [Operations] is where risk management acts a control. So this is where risk management would give us insight into what we already do. This is where you make explicit what you already do, which is interesting, cos that's what you were talking about the other day about institutional learning. But, then, so much of what I seem to do here is try to capture what the business already does, and I know there's a driver for that in terms of governance, but what value does it bring? Part of it could be, you could argue that it gives you a better understanding of why you do things.

RD: Well yes, because if you limited the use of risk analysis only to here [pointing to Business Projects], then you're effectively only concerned with capital expenditure, with significant decisions on a new capability, which in this organisation is mostly a physical asset being constructed or put in the ground. You would miss all those places where day-to-day activities, systems, and processes were risk controls. You would have no idea as to why those things existed. So, for example, the health and safety procedures, induction processes, or procedures for managing financial transactions.

CRM: Yeah ok, so there's value in that, in that it will hopefully encourage people to think about what they're doing. And if they think about it then perhaps they can change it for the better. So that's what we're saying, by making decisions, any decisions, based on risk data, we will make better decisions. If we compare what we're doing against risk data it should give us a basis for saying whether or not we could do things better. And that's what RCM is doing isn't it? They've collected some data

about the existing maintenance processes and used that data as a basis to alter the maintenance programmes, to do them differently. Now each of those decisions is not a significant decision, but where it's significant is because you do so much maintenance.

RD: Yes. I think that's what is called reflective practice, getting people to think about what they're doing.

CRM: Ah, I like that term. But how do I sell it? I mean, they're just going to say 'I know what I'm doing, don't come and look at my group'. Hmm, yeah ok. Maybe the fundamental thing is that they need to appreciate that by having a look at what they're doing there's value which comes out of the point that studying what you're doing gives you the means to do things better.



## Dialogue No. 12 – On the value of the risk framework in a mature organisation

**Source:** Transcript 23 November 2007 (CRM:18)

**Participants:** Corporate Risk Manager (CRM)  
Richard Donnelly (RD) [*Researcher*]

---

CRM: Something else that dawned on me is, in essence, what I've done is to walk around the business, have a talk to people, and worked out how we do things, and what I'm finding is I'm constructing my framework in a way to take advantage, or to capture work that's already done in the business; which is ironic then, isn't it, because you could ask "how much am I adding to the business?"

RD: Yes, I guess you could ask that. But you could also look at it from an empirical point of view. What's already going on, the existence of those metrics and the systems to measure those things would suggest that they are important to the organisation. So their very existence verifies what you've come up with.

CRM: Yes, that's right the compliance management system is indicative of the fact that statutory compliance is important to the organisation.

RD: If these things didn't exist then you would have to ask the question, have you got it wrong? If you had taken a different path, dissected the company differently, and created a framework that required entirely new calculative capabilities you might then face the objection of "why do we need to develop new capabilities when we already have these?"

CRM: Yes that's true, but part of me wonders about the fact that I would call Watercare a mature organisation, in the sense that the company, in whatever form, has been around for the better part of half a century. If you look at it from an organisational knowledge perspective, where the business develops and retains knowledge about how it goes about its business, then presumably, over time, the organisation has developed better and better ways of doing things. Now Watercare's been around a long time, so when I look at the business and what it does I believe, and I have believed right from the beginning, that they're pretty good at managing the common risks associated with the business. Sure, there are probably little holes in amongst what they do, where the company might be exposed, but in terms of the core stuff I think they're probably pretty good at it. But the irony is that I think the risk framework is well behind in the sense that it's not representing very well what the business already does. So I think the initial challenge is develop a risk framework that accurately reflects what the business currently does, and then once you have

achieved that, your risk management system should ideally be capable of providing insight into what the business doesn't do well. And so as I'm developing the risk framework it doesn't surprise me that while I'm listening to the people on the coal face, looking at the strategic objectives, and trying to marry them, that I'm finding that a lot of things are coming into alignment, because I believe that Watercare, as a mature organisation, has become good at what they do. And I wonder if you went into an immature organisation that you could set up a relatively crude risk framework and very quickly pick holes or areas for improvement.

RD: So you're thinking that potentially risk management adds more value in an immature or young organisation because, by definition, such an organisation is still learning how to do what it does, while a mature organisation like Watercare already has well developed processes?

CRM: Yes just like a graduate starting a job. Initially you should be able to point out lots of ways they were doing things wrong and could improve, but if you look at someone who's fifty-nine years old and been doing the same thing forever, there's very little I would expect to be able to tell them. And for me that really calls into question the value of risk management. I mean, I think there are two primary drivers for risk management. Certainly, from the Board's perspective, I believe the reason the whole risk management function was setup was this concept of protecting the organisation. So from a corporate governance point of view they seem to want to have someone monitoring the company through a formal system. My objective there, from a governance point of view, is to fairly reflect what the business does, which I actually think is more concerned with analysing the business, finding out where there are holes and then also checking that the I's and T's are dotted and crossed. To be fair, at the moment, I think the risk management function does not fairly reflect what the business does, not to say that our reporting to the Board doesn't, the two could be different things. But from a corporate governance perspective I will feel better about things when I feel that our data is a more realistic representation of our understanding of our risks. But it seems like there is enormous overlap in that function with the internal auditor because the internal auditor is meant check that the I's and T's are dotted. Now you could say the purpose of the risk management function is to analyse where there are holes, but this is a mature business so what's the chance that I'm going to find big holes? In which case, what's the value? Is it worth spending my salary every year plus the distraction I cause to other staff just to find what are likely to be little holes in the organisation?

RD: Put like that, it sounds like a lot of effort for little gain. But presumably that's where the value is from a corporate governance perspective. That the governance function of risk management is fulfilled by this process of making explicit what the organisation already does, having a formal procedure to verify that the organisation knows what it's doing.

CRM: Well that's the irony of it. Basically I'm trying to put mathematics onto reality, to draw a picture using mathematics that captures and conveys reality. So yes, in that sense, even though the risk management framework is, in many ways, simply trying to catch up with the organisation's practices, presumably there is value in the sense that that is precisely the role of risk management from a governance perspective, to make sure that the Board is getting a realistic representation of the company through the risk analysis and data. Actually that brings me to another reason for risk management, although I'm not completely sure whether it fits under the governance banner, and that is that risk management improves the defensibility of decision making. At Watercare I think there's a practical need for defensibility driven by the political environment in Auckland. I mean, rates are continually going up, there is already pressure on Metrowater because of the price of water, and we know from our own in-house analysis that there are going to be large infrastructure investments required in the future, so it's foreseeable that Watercare's [Asset Management Plan], or requests for funding through its AMP are going to come under heavier and heavier scrutiny in the future. So I think the practical need there, and the value that I see my work adding to the business is with helping to formulate the AMP. If I can get the risk data to a point where they can analytically calculate the AMP scenarios then presumably the AMP will be a lot more defensible than at present where the decisions on which projects to cut are somewhat more subjective. So at a higher level I tend to think there's some kind of defensibility value there as well, which you could paint as providing a more objective basis for decision making.

RD: Ok, so that's interesting. In both cases there you highlighted the role that risk management plays in providing objectivity, so a realistic representation of the business, and objective data to support the AMP, both of which come from the analytics, the formal analysis.

CRM: Yes, I guess you could say I'm analysing the business from the risk management standpoint. Actually no, I'm not actually analysing the business. What I'm trying to do is to set up a framework so that people in different areas of the business will enter their knowledge and then that framework will process knowledge from different areas of the business and then represent, draw a picture of the business. So it's meant to be, in theory, I guess, an objective way of looking at the company, but more importantly it's meant to be an enterprise wide view. I'm meant to be looking at everything you could compare across the company. The bottom line is that it's supposed to compare everything within the company on common grounds, and those common grounds are the overarching corporate objectives.

RD: Ok, so are those the only ways in which risk management adds value to the organisation.

CRM: Well, I'm still struggling with that to be honest. I notice that in the risk management journals...

RD: Ah, which journals?

CRM: There are two that I primarily see. One's called Strategic Risk from the UK, and the other comes from the Institute of Risk Management which is a very light weight document. I believe the articles in both are generally published either by consultants or local government practitioners.

RD: Thanks. Sorry, carry on...

CRM: Well I notice that there is a big drive to sell the commercial value of risk management, and I guess I'm a little sceptical about that. I mean from the governance point of view, and even in terms of the defensibility thing, risk management seems to be really an administrative function. It's about justifying what the company already does. But then, the theory for risk management says that you should be able to add value in the sense that the risk analysis should provide additional insight. But I'm not sure about that here yet because I feel that I'm behind what the company currently does, and I'm still trying to catch up. As I said, I think that Watercare is a fairly mature organisation and I think you're going to have to work at a very high level of detail to provide additional insight into what they do. So I think it's more challenging in a mature, well developed, well experienced organisation, for the risk management function to provide additional insight because it involves having a very complex or very detailed look at things in that kind of organisation. What you could perhaps argue is that, in an environment where there is greater and greater pressure to increase your output, you could argue that what the risk management function does, and I'm sure that older, more experienced engineers would argue, that there's less attention paid to detail now. You often hear that when you talk to older engineers. They say "we used to make sure that the As-built drawings were done." The quality of workmanship is shrinking, and part of me wonders if the growing pressure for improved productivity comes at the expense of reliability at some level.

RD: Sorry can you explain that?

CRM: Right, the point that I was trying to make is that it's probably going to be difficult to analyse what Watercare does and find big holes or big risk exposures because I believe the company is what I would call a mature organisation. However, what you might find is that by forcing the engineers to sit down and think about the various sources of risk and to think about the actual business systems that are in place to mitigate different areas of risk, what you might make them realise is why they do things. Some how that ties back to the concept I was talking about before about dotting the I's and crossing the T's. I think that in the past they were better at it, where now we seem to work more on the eighty-twenty rule and I think that's probably got reliability implications. Part of me wonders that if you sit down and analyse the business and why it has these practices what you might realise is why. So if you sit down and analyse the business and you interpret existing business



practices in terms of the role that they might play in terms of risk management, does that make sense...

RD: Yes, I'm following...

CRM: ... so if you're interpreting a business practice in terms of a risk management system you might recognise then, in some ways you might better understand the purpose of that business process. And if you understand the purpose of that business process, then you might appreciate the relative importance of different business processes, and if you understand the relative importance of different business processes then you might change the way you behave, I guess. As opposed to sending a graduate into a room and saying "do those calculations, send out those documents, draw those drawings", and he does them and when he's short for time, he cuts corners wherever he can. If, after a year of doing that kind of work, you said to that graduate "look at what you do and why you do it in terms of risk management", they might say "actually, producing the drawings at the end is the most important thing, so if I'm going to cut corners I'll cut them at another point", I guess. So, I'm not sure, but in terms of a mature organisation you're not necessarily providing new insight. In some ways all you're doing is providing justification or almost prioritisation of what they're already doing. But the bottom line in some ways is that concept that you've mentioned to me before, that the risk analysis process is almost making them reflect on what they do and why they do it, and in some ways the risk data is the relative importance of different things.

RD: Yes, that makes a lot of sense because risk management is a type of process itself, a form of inquiry for thinking about other processes. And I think you're right. In a mature organisation that's where the value comes from. It's about reflection and learning, getting people to think about what they're going. Whereas, in contrast, the example you had at the beginning, risk management in an immature organisation, or perhaps in an organisation in a complex and rapidly changing environment, is very much about identifying events or opportunities, actively managing the business process in relation to the environment. Now that seems to suggest an important link between the context of the organisation and the nature of risk management in that context. I wonder, given the drive you mentioned in the literature to promote the commercial value of risk management, if that is giving rise to misconceptions in the Boardroom or executive level as to what risk management should be delivering in organisation, and especially in mature organisations like this one?

CRM: Yes, it's not black and white is it? It's not simple. That would be an interesting exercise to put the managers on the spot and ask them exactly what they think risk management should be delivering. But let me just come back to a point you made there. I said for an immature organisation you might provide an insight into areas where there are big holes and by converse in mature organisation's all you might provide is a better appreciation of why you do things, the priority of why you do

things. Now I'm not sure that the two are different. They're just at different ends of the spectrum, in the sense that for an immature organisation you might say that you need to develop a new business process because that capability is missing, where for a mature organisation the priority might be on further developing existing processes. I think they're just different ends of a spectrum.

RD: That's interesting. You've rung a bell there. There's a similarity I think with the knowledge management concepts of evolving one's knowledge from "Know that" to "Know how" to "Know why"...

CRM: Yeah, so we know why we do it...

RD: Yes that's right. The concept being that we start off with "Know that" knowledge, or knowing what we should do. We then advance to "Know how" knowledge, where we know how to do something, and then advance to "Know why" knowledge, which is where we understand why we do something. The interesting thing is that then links, at least, in theory, risk management to innovation, since product and process innovation only becomes possible once you reach the "Know why" state.

CRM: Yes ok. There's an analogy that keeps rolling round in the back of my head, which I picked up at some sort of leadership training and management course that I went to in the past. Inevitably the focus of those courses is self-awareness and they had this phrase that I thought was incredibly corny. It was very American. They said that as a manager you have responsibility. And it was this phrase, you have responsibility, and they hammered it throughout this course. What they were saying was analyse that phrase, you have response-ability, which is, and their whole idea was, that as a manager you have the ability to choose how you respond, you need to be aware of that, you need to have a wider awareness, you need to have self-awareness. And so I think that aligns very well with what you've talked about before, this idea of reflective thinking. When something comes up do you just react or do you think about how you're going to respond, what action you should take? So there are good parallels in terms of personal development and organisational development. That's all that I seem to be doing is trying to generate self-awareness for the organisation.

## Dialogue No. 13 – Trying to explain the role

**Source:** Transcript 3 March 2008 (CRM:19)

**Participants:** Corporate Risk Manager (CRM)  
Richard Donnelly (RD) [*Researcher*]

---

CRM: Over the last couple of weeks I've been thinking about things that I want to talk about at the conference. Some of the things that have popped up are things we have discussed in the past, like what are the objectives for risk management, corporate governance, insurance, the reliability of achieving objectives, the role of risk data within the organisation, the level of detail being driven by the AMP here at Watercare. But the first thing I wanted to cover was to ask "what is Enterprise Risk Management?" So I've been looking in the risk management standards and there's a bunch of key words, so integration, business processes being embedded, being enterprise wide, full depth, full breadth of the organisation, holistic, all sources of risk and continuous, forward looking. But I wanted to talk about the fact that there's a method gap, in terms of how to go from these prescribed standards to what is essentially a vision, a culture. So I started thinking about some of the things we've seen around here. So the risk register for example. At the moment I own it, the data sits with me, it's a detached thing. We've realised now, although its current best practice in New Zealand, its an illustration of risk management being detached from the day-to-day business. It's not integrated. So I think this process of getting someone to come in, buying them in as a risk manager, getting them to set up a risk register, and then periodically reviewing it by getting staff to come into their office once every six months, that's a bolt-on solution. That's one observation. Another is to do with this concept of being holistic, looking at all sources of risk affecting the business performance. Well, you can measure the performance of the business in a variety of terms, how much money it makes, or how many apples it makes fall off a tree. Point is, you decide what measures you want to use, and the process that we went through in terms of setting up the risk framework was really a process for working out what matters in terms of performance, and we came up with five measures of performance that we were going to use.

RD: That's a good phrase, what matters.

CRM: Yeah, you've got a raft of possible ways you could measure organisational performance. You end up going through a process where you're trying to boil it down to a handful of those that matter. And then there's this thing being enterprise-wide, which means you've got to have risk management across the organisation but also throughout the full depth of the organisation. Being across the organisation puts the focus on how the organisation as a whole performs, not how the chapters

of the organisation perform individually, so it focuses everything on a key driver, a common driver, which requires you to think about what the principal objectives are for the organisation. But this other part, about being full depth, I think relates to what we have talked about before, the marrying top down with bottom up, because you need risk management activity at lower levels to be correctly aligned and supporting decision making at the top. So I don't really know how to convey all of these things yet, but I did realise that a lot of the issues that you and I have previously discussed could be easily be bracketed in some of these terms here.

RD: Well, I had some thoughts as you were talking. You brought up the concept of the integration, about risk management being embedded into all aspects of organisational decision making, and you talked about the risk register being a bolt-on solution. If you think about the role that the risk register plays in knowledge production, something I have realised is that the risk register merely captures the end product of that process. So there are various processes in the organisation that generate knowledge for decision making, investigations, analysis, modelling, they all generate data that goes into formulating a project or decision to spend money. The risk register captures a statement of risk which is the output of that knowledge production process, but such statements can't be constructed until that process is nearly complete. So where the risk register simply sits as a recording device at the end of that process, it's simply a bolt-on. But this notion of embeddedness could be seen as actually managing that knowledge production process...

CRM: That's right, the generation of it...

RD: Right. And this other concept you were talking about, being enterprise wide, being full depth. You mentioned marrying top to bottom and I had two thoughts. The first was that the link between risk management and roles and responsibilities is relevant there, so the task of aligning responsibilities for risk with those who can actually control it. The second was when you mentioned that risk management activity at the lower levels needs to support what's happening at the top, I thought brings in the idea that there needs to be a translation of information from the bottom of the organisation to the top.

CRM: Yes that's right, translation was a key word we used to throw about.

RD: Yes, we had talked about how the risk framework acts as a mechanism for translation.

CRM: It's funny you use that phrase, because just today I was thinking about my approach here. I mean a natural approach is to employ someone, get them to establish a risk register, bolt it on in effect. But really what I've done is to have a look at our existing business practices, and the fact that we have these existing features here, so for example the health and safety management system, the compliance management system, and so on, these are all evidence of what could be interpreted as risk management activities. In essence, I realised that if I was going to be efficient in

what I was doing, I shouldn't try to re-engineer the business but to integrate risk into it. It's about interpreting what these things are under a title of risk, if you like, finding a common basis for assessment across the business, finding a way to compare apples with oranges, that's this whole thing about risk being a common basis on which to make a comparison. So some one could come in cold, set up a new system and say "I'm going to apply my system to your business and compare things in your business based on my system". Or you can come in and look at the business and say what here would give me a good basis for common comparison. And, really, that basis is the corporate objectives, and what risk management tries to do is to compare everything that goes in within the organisation on the basis of how it could impact on the achievement of those objectives. So I think it's about interpreting the existing business practices from the risk management perspective, but coming up with that framework is really part art, part science. You need to be able to look at a business and understand it, and at the same time you need to be able to look at your objectives and deconstruct them, and then try to marry the two together. You have some paper-based objectives, and then you have real business practices and what you're trying to do is to bring these two together to achieve an efficient translation of information from the bottom to the top. Hmm, I'm just recognising, talking through these concepts, that there's too much. I won't be able to explain all of this.

RD: Yes, you're covering a lot of evolved understanding there.

CRM: I had hoped to talk about some basic things, give some commentary on Watercare's experience, but if I just get up give three examples of challenges that we've faced, and then walk away, I haven't really added much. Sure I can make the point that it's complicated, it's hard, but I've not really added much, I haven't really pulled anything together. Granted, there's no simple solution, which is I guess where this bit about it being as much art as science comes it, but that's pretty vague.

RD: Well, a thought, on that last slide you were talking about interpreting the business from a risk management perspective, which emphasises that a lot of risk management activities and processes already exist in the business and there's no need to completely reinvent the wheel.

CRM: Yeah, which I think is a very important point to make. Risk management isn't something that you come in and set up completely from scratch. Most businesses are already doing things. Good practice, good management is already risk management. Certainly, when I started to try to explain these problems, or even just to put these issues into words, I realised that I was being drawn back to the fundamental question of what are you doing here? Fundamentally, what are you trying to achieve? What do you need data for? What is it that you're supposed to do? What role does the data play? Because the risk management process is basically the process for decision-making, it closely resembles the rational decision making model. But Enterprise Risk Management is also talked about as a culture.

So this is the question: how do you go from what is essentially a business process definition to a culture? I don't know.

RD: Right, so how do you take an abstract process concept and embed it in what people do everyday?

CRM: Yeah, so that it underpins what they're doing. And what we're saying is that there are already elements of this process that we already do. Well of course, there must be. People are making decisions all the time about what's the best course of action for the organisation, weighing up the options. So even if they're doing it sub-consciously, like what you observed in that Hunua No 4 workshop, where they were jumping to the answers without all the processing in between.

RD: Yes they intuitively knew the answers, and so that's what the discussion quickly focussed on. Hmm, I wonder if it's a case of turning it around. So there is this abstract process that is published in the standards and guidelines, but its abstract, it's a generic vision, which is a term you used. And you asked the question of how to move from that to a culture. But your experience has shown that a lot of risk management is actually business as usual, it's already being performed, and a significant part of what you have been doing is capturing explicitly what is already happening implicitly in the organisation...

CRM: Yes, that's right, reinterpreting what's already done. These are phrases that we used all the time.

RD: Yes they keep coming up. So maybe it's not a case of asking how do you take the process to the organisation, but rather of asking how do you take what already exists, what's already happening, and make it more explicitly resemble the generic risk management process?

CRM: That's it, yeah. Which is what I was basically saying on the next slide. This is what's happening. So how do I put it in Enterprise Risk Management terms? Actually that's not a bad way of looking at it. So, in a way, re-representing current business practices in terms of the ERM concept...

RD: Yeah, it's almost represented by the terms you used on the previous slide. Culture is an implicit term. Culture is implicit in an organisation, it's how you understand the way things are done in a company, but it's not necessarily the explicit rules and regulations...

CRM: Ah, this is getting interesting now. So culture's implicit, process is explicit. So rational process versus culture, which is intuitive, isn't it? I'll show you a phrase that I just dug up. I have this quote, from a song, "so you spend time trying to get your hands on the holy grail", and what I was trying to say here is it's a short song, a short presentation, "it's a hell of a story", I'm presenting material that only barely scratches the surface of some of the issues that I've been grappling with of effectively achieving the "nirvana" state of ERM. Now I never knew really what I was

thinking there, but when I looked up what “nirvana” means, in Buddhism nirvana is a state of peace and enlightenment that involves being unaware of one’s self, and the phrase that was used was being “unconsciously competent”. Another definition was that it’s the Buddhist term for self-realisation, the transcendence of suffering. Well, going through the risk management process is a process of suffering, it’s an administrative process, and what we hope to get to is the transcendence of that suffering where you have complete self-realisation and you’re unconsciously competent.

RD: Hmm, that’s interesting, there are multiple parallels there. We’ve got “culture vs process”, that’s one way of putting it. Another phrase we used was “implicit vs explicit”. And still another way of looking at it, taking an analogy from the knowledge management literature, is “tacit knowledge versus explicit knowledge”. So, what you just talked about there, the process of suffering leading to the nirvana state of self-realisation, that seems to describe the process of organisational learning. So you tacitly understand what you’re doing, its intuitive. Then you go through an explicit process of documenting that knowledge and reflecting on it, so that’s the risk analysis process, the process of suffering, which then leads you to further tacit knowledge. It’s a cycle.

CRM: These are really good parallels, the state of enlightenment, and being unaware of ones self, you’re not conscious of risk management, you’re just doing it, it’s a good practice. That’s really interesting.





## Dialogue No. 14 – Reflections on a year in the role

**Source:** Transcript 23 May 2008 (CRM:20)

**Participants:** Corporate Risk Manager (CRM)  
Richard Donnelly (RD) [*Researcher*]

---

CRM: Thinking back to the beginning of last year, you were interested in data quality and you felt that risk management should be supporting decision making by providing good quality data, and that framed your approach I think. What I'm trying to get at now is to understand where you're at now in terms of your thinking.

CRM: When I first got here, I recall using the term "defensibility" a lot. I was talking a lot about the risk data. I thought that the risk management system had a huge role to play in terms of providing the organisation with defensibility for the way it made its decisions. So the data should justify the decision making and back then I wrote that generating data was one of my fundamental roles, to justify or provide a basis for decision making. And data quality is still something that seems to be very important to me. Certainly, from a corporate governance perspective, I see that my role is to draw a picture of the organisation, to illustrate where we are perhaps most vulnerable, to show the threats to performance. But I guess I'm seeing that the data itself is less important in a way, but it's a difficult point to explain. I see my role differently now. I see my role now as primarily, internally, to influence thinking, to influence the way that people think about the organisation, the way that they see the organisation internally, how they see their internal customers, and how they might view their relationships with other people in the business, so that they have better clarity about what it is they're required to do and why they are required to do that. So, at this point in time, I think my role is primarily concerned with influencing how people understand what they're doing and why they're doing it, which is the cultural change that I'm trying to achieve.

RD: That's quite a turn around. How did you start to realise that it was more about the thinking than about the data?

CRM: Well it was a challenge. I think my background as an engineer certainly influenced the way I read the risk management standards. I guess someone from a creative background might have read them differently, but being an engineer, when I read the standards I saw a defined step by step process and it looked to me like it was all about generating numbers. I thought they implied an analytical methodology behind risk management.

RD: I know what you're talking about. That's precisely how I interpreted the standards, and I think my own training as an engineer had a lot to do with that.

CRM: Right. So I guess I came into this role with a certain expectation that there would be some purity in the analytical approach, that you could generate numbers nicely. But what I realised is that there's no perfect analytical method. Some are better than others in terms of generating objective data but all of them are subjective and you can measure things a multitude of different ways. And so I realised that you don't have to be analytical. The conclusion you come to is that risk management may not necessarily fulfil its roles by being analytical. Sure, there is an analysis role there, but that analysis doesn't have to be quantifiable. Basically it's about identifying areas that warrant improvement or attention. And that's, as you and I have discussed, is about getting people to reflect on what it is that they do and how they do it. In some ways my role is about providing a more thorough or clearer context for decision making by other people, so influencing people about the bigger picture. So, for example, I spoke to Operations group and pointed out that there's a whole lot of operational risks in the register, and there's a perception in the organisation that risk management is an operational issue. And then I said "but did you realise that the people who control the risk environment are actually the Asset Management group?" And that changed how a few people were thinking about the organisation.

RD: Specifically?

CRM: Well, [the Networks manager] came up to me after the presentation and said that he had actually never thought about it that way. He had always thought about it the other way, that Operations were the upstream party and that Asset Management were the downstream party. And none of that came about through data. That realisation came about through discussion. So my role is to influence thinking, and influencing thinking may not actually require data generation. The problem there though is that I think the data is expected. So people, both inside and outside the company, expect to see risk data if you're doing risk management. They expect to see a risk register with measures of things.

RD: Is that something you get a lot? Why do you perceive that expectation?

CRM: Actually that's speculation. But I would expect that if [the external auditor] came in here next year and said "show me your risk information", and all I had was a list of issues with no explicit measurement framework, I think they would criticise the method because it's not analytical. Where in actual fact I might be more effective by travelling around the organisation and working with people to get them to think about what it is that they do, why they do it, who their customers are, what they expect, identifying issues to address, and working on improving things that way, rather than by spending all this time trying to be analytical in an environment where there is no precise analytical solution. Now people would debate that with me, perhaps because they don't understand everything that we have discussed. But I've realised that, in some ways, the data is almost pointless because you always compare the analytical data to your gut feeling, you determine whether or not it's right by whether or not it matches your feeling. The classic one there was the project

risk register. I found that one of the project managers had inserted a worksheet called "Sanity Check". She would put all the data into the register and do the assessments, but then she would sort them to see if say the top ten risks looked alright. If they didn't she'd go back and adjust her scoring. So that made me think, what's the purpose of having data if all you're trying to do is to reinforce what you already know?

RD: Yes, there was a similar episode when you were talking to [the Principal Water Planner] about developing a tool for the analytical formulation of the AMP, and when you asked him if he thought such a tool would be helpful, he responded that it was a question of calibrating the analytical data against the outcomes they currently have. I think that's the same thing. The objective data gets calibrated by the subjective judgement.

CRM: Yes, and I think there is an element of reflecting on what you're doing there. But I suspect where more value exists is when I go down to the project managers and I say "do you understand the purpose of the contract, do you understand the scope of insurance cover and when it's appropriate, when it's inappropriate?" That's got nothing to do with data. I mean, we could spend ages generating data for me to show that the contract and the insurance are the two primary controls there, but we already know that. So what's probably more effective is for me to provide training to the project managers on these critical elements of practice. It's just good practice for them to understand the contract well. And another thing, I don't get positive feedback from the business about the data, like "Oh yeah that data's really good". At the end of a meeting, people say "that was really interesting", or "that was a really good meeting". But it's only because people have been around a table, talking and debating, and it's changed the way they've been thinking, or they've made some progress or got agreement on something. So I think that the organisation really does get value out of reflecting on what it does, and out of having someone like me critiquing the organisation and talking to people about it, and trying to get them to think about what they do and how they do it, or getting them to view the organisation in a different way to how they normally see it in their day to day business, to stand back and look at the big picture. And you don't need data to do that.

RD: Thinking back to your earlier focus on data quality, it almost seems like your point of view has shifted 180 degrees.

CRM: Yeah, it's definitely a 180 degree turn, from saying that I've got to generate good quality data to saying that I could do this job by not generating data, that all the job requires to be done efficiently is to have someone who's insightful. You need someone who can look at systems, patterns, interactions, and who can cut to the chase. But the irony is that there's still a part of me that says "but data quality's everything, data quality's everything", and that must be the engineer. I suspect if you read this transcript against the first transcripts from way back in May or June last year, you will see back then that I thought things were so much more decisive, so

much more analytical. As you say, I was talking about data quality for objectivity, and now I'm at the opposite end of the spectrum saying nothing's simple, it's all complicated, there's no right answer, so I guess you could say it's a complete paradigm shift. But the irony is I feel like I'm better focussed on the core purpose of risk management, which is really concerned with making the organisation less vulnerable, a process of continuous improvement. I think now I could do it a lot better by just going out, talking to people, looking at the business, cutting it and slicing it a hundred different ways, until we see things that drop out as important. Although that still leaves you with the problem of how do you report to the Board on an objective basis? How do you provide the Board with an indication that things are well managed if someone is just floating around, trying to facilitate development where they think it's needed? I guess that's why you need to have data in the end.

RD: Well that then almost takes you back full circle, in the sense that you have this person floating around who is driving change, but then he or she has to defend what they're doing.

CRM: Yeah that's right, defensibility of their role. The irony there is, I'm starting to see that the data, the process of generating a lot of detailed data, is probably, not a waste of time, but a huge inefficiency, but then, without it, how do you justify that you're focussing on the right issues? It would come down entirely to individual opinion. You would rely on the risk manager to be someone who was very insightful, someone who had a comprehensive view of the business, to be focussing on the areas that really need attention. It would rely on my effectiveness, on my being 100% on the money with my thinking. So, for instance, fundamentally I've started to recognise that the fundamental weakness in the business structure is actually communication between the two main functional groups, Operations and Asset Management. You've really got to have Operations telling Asset Management what the situation is, and then you've got to have Asset Management understanding that they heavily influence, if not control the operational environment. If they don't understand that then you have all sorts of problems. Basically everything flows from that fundamental need. And so I agree with something that you mentioned, that the risk manager role is an ideal position to put someone before they move into a more function-specific role because it forces them to take a detailed look at the company, to see how it interacts, what's important, what's not, what are the fundamental roles and relationships, and what are the risks. I think that is a clear part of my role. So, for example, I noticed that maintenance planning and condition assessment are located in the operations group, and yet those are primary asset planning functions. So I think there's something wrong there, and I've raised it with some of the General Managers, who've now realised that there's an inconsistency there. So there's an example of where I might be adding value through the way that I've viewed the company, and then drawn other people's attention to the problem. And I didn't need a lot of detailed data to do that. So it seems silly, to me the most inefficient thing about my salary is that we need to see a risk register, we need to see data, and yet

really, the data does nothing, or very little I think. All I've been trying to do is to generate data to illustrate what I think, intuitively, is going on. And I'm hoping that the data verifies that the business is working in the right areas.

RD: So, if you had to choose, what would you say is the key understanding that you will take away from your experience here?

CRM: I think the key understanding that I will take away from this is that, in essence, risk management is a polarising lens for viewing the business, something is either a risk or a control, and it is my role to interpret the organisation in terms of threats and controls. So I came into Watercare thinking about risk management as a bolt on process, but I've come to realise now that risk management is not some kind of new process that the organisation is going to start doing. It's simply a way of looking at the organisation, at what the business already does. So, in a way, risk management doesn't exist, but the business already does risk management, and you just have to look at in those terms. That's something that's dawned on me over the last couple of months, that it's about how I view the organisation, because I've spent so much time looking at the risk framework. In fact it's only in the last month that I've moved on to the second step of the risk management process, identifying risks. And I've realised that identifying and describing risks is actually not simple. It's a question of how much detail is appropriate. So I started with the two principal objectives of the organisation and I deconstructed them into what turned out to be seven risks, or, put another way, seven questions each asking how we could fail to achieve something. And two of those I've now deconstructed into a lot more detail, and through that process I've gone from seven risks at this high level to what might be 50 risks at the next level down, and then we could go down into further detail, and I could describe what might be 300 or 800 risks at the next level. So each risk can be further deconstructed into lower-level risks, or smaller and smaller sources, and it's a question of where is the appropriate level of detail? And that's where, and I've discussed this with you before, I think that the process I'm going through of describing and representing the organisation's risks and controls, of making them explicit, needs to marry with the operational needs. So we've already got a health and safety hazard register, a compliance management system, public health management plans, and I would like to separate out the asset risks into a separate register, as well as having a high level register for governance reporting. But the point is each one just describes risks in different levels of detail which marry up with the specific needs of the business at that level. So, for instance, in the register I've standardised the description of controls so that I can interrogate the data to find out information relevant to specific risks and controls. But it dawned on me that [the Internal Auditor] is also interested in controls, so where I define what the controls are, why don't I also define the assurance history for each control. Say one control is monitoring compliance with the treasury policy. Well in 2004 [the external auditor] came in and did an audit on that point, in 2006 [the Internal Auditor] did an audit on that point, and all of a sudden you have your link between risk and assurance, the

assurance history. And that made me realise that the way I describe those controls isn't absolute, it needs to suit the needs of [the Internal Auditor], I've got to come up with a way of describing controls so that he can provide assurance over them. But it's not a simple process, there's a bit of an art to it, it's not pure science, it's not well defined black and white.

RD: Can you elaborate at all on how it is part art, part science?

CRM: Well, when I used this phrase more art than science, it reminded me about something we had previously discussed, tacit knowledge versus explicit knowledge. Now tacit knowledge is something you understand but you can't necessarily write it down. But to me tacit knowledge is probably the determining factor in how effective risk management is rather than explicit knowledge, and I'm recognising that I can't write down a lot of what I'm doing here. It's very difficult to sit here and tell you how I came up with them you know. And so it's made me realise that's the art side of risk management. I mean, my framework is explicit knowledge, an explicit representation of the business. And sure, I've gone through a process of analysis to generate it, but it's relatively crude and I don't believe it is as influential, in terms of influencing behaviour, as looking at the organisation, recognising where there are problems, seeing how things interact, and then going to someone and saying "did you think about it this way?" In essence my job is to say to the business "ok, we've done some risk analysis, have a look at the results of the analysis", and what I think will happen is people will either look at the data and say "that's good", or they'll look at it and say "aw, you've missed a whole lot of stuff there". Now, it's not that the analysis may have missed anything, but maybe the things that they're talking about are identified and in ways that they're not used to thinking about. So the art is very much, I think, about how you marry the analysis with the way the organisation thinks. I mean, how often have I stood up to present a risk analysis, and someone asks "what about this or that?" and what they're talking about is there, and I have to point out that they're concerns are just represented differently. After a while they shut up, but you can see that the way I might approach the analysis doesn't necessarily marry with the way everyone's way of thinking. So I think it comes back to this thing that risk management is not necessarily a new process. It's just a way of viewing the business, like a polarising lens so you can view a company in terms of threats and controls. Risk management's a way of looking at a company, not something new that a company does.

# Appendices

*Refer to included CD-ROM (print version only)*





# Appendix I

## The research and practice context

The following sections describe the institutional and governance context of Watercare Services Ltd (Watercare or WSL), certain features of the company's internal structure and functions, and the risk assessment framework that was in use within the company prior to 2007. The information presented in this Appendix refers to the company as it existed prior to the restructuring of Auckland governance in 2009 and 2010.

### **The corporate context**

The Greater Auckland Region encompasses a land area of 5,020 sq km, extending 120 km north to south, and 60 km from west coast to east coast at the widest point (McDonald and Patterson 2003). At the time this research was undertaken the population of the region was approx. 1.4 million (Statistics New Zealand 2007; estimate based on 2006 census), distributed across seven cities and districts: Rodney District (pop 95,000), North Shore City (pop 220,000), Waitakere City (pop 198,000), Auckland City (pop 433,000), Manukau City (pop 355,000), Papakura District (pop 48,000), and Franklin District (pop 62,000).

Water supply and wastewater services were provided to the main urban populations of the region by Local Network Operators (LNOs). There was one LNO for each city or district, with the exception of Franklin District. In Franklin District, and in rural and isolated urban areas of the other cities and districts, these services were provided either by standalone systems owned and operated by the local government authority, or by individual landowners (i.e. rain water collection for supply and septic tank disposal of wastewater). The LNOs

developed, operated, and maintained the local area water supply and wastewater infrastructure on behalf of the local authorities, under a mixture of governance arrangements: three of the LNOs were business units of the respective local authority, two were stand alone profit generating businesses wholly owned by the respective local authorities, and one was an independent international water service company (United Water) operating under a long-term contract to the local authority. Stormwater services are wholly provided by the local authority in each city or district.

Watercare Services Ltd was the bulk supplier of water supply and wastewater services for the Greater Auckland Region (maps of the geographical distribution of Watercare's infrastructure can typically be found in the company's annual reports and annual asset management plans). This means that Watercare sourced and treated raw fresh water, distributed the treated drinking water to the local reticulation networks within the region, collected wastewater from the local collection networks, and transported, treated, and disposed of that wastewater. The relationship between Watercare and LNOs was thus one of wholesaler to retailer. Watercare, the LNOs, and the local government authorities, together constituted the Auckland Water Industry. In total, this industry supplied over 141 million cubic metres of water, and collected, treated, and disposed of over 133 million cubic metres of wastewater per annum (Auckland Water Group 2007). The industry had an annual turnover of over NZ\$600 million, operated over 22,100 km of water, wastewater, and stormwater pipelines, employed over 800 full time personnel, and was supported by around 1,000 independent contractors and consultants (Auckland Water Group 2007).

## **Infrastructure, services and performance standards**

Watercare operated, maintained and developed the raw water sources and headworks, water treatment plants and bulk distribution system, the bulk sewer network, and the Mangere Wastewater Treatment Plant (see Table I.1 for a summary of Watercare's infrastructure). On the water supply side the company operated a total of twelve raw water sources and six water treatment plants, of which the Ardmore treatment plant was the largest, with a total treatment capacity of 333,000 m<sup>3</sup> per day. The geographical area serviced by each LNO was divided into supply zones and Watercare supplied treated water to designated points within those zones called Bulk Supply Points (BSP). A Bulk Supply Point (BSP) constituted the boundary between Watercare's network and the network controlled by the LNO. It was also the point of sale for treated water from Watercare's network, and the point at which key levels of service were defined and measured. Watercare operated a total 444 km of bulk water

mains, ranging in size from 300mm to 1900mm diameter (Watercare Services Ltd. 2007b).

Levels of service for water supply were governed by the Ministry of Health's New Zealand Drinking Water Standards (Ministry of Health 2005) and drinking water grading (Ministry of Health 2003), and the customer contracts between Watercare and the LNOs. The DWSNZ specified maximum acceptable values for the microbial, chemical and radiological determinands of public health significance in drinking-water, provided compliance criteria and procedures for verifying the water supply was not exceeding these values, and specified the actions to be followed when a transgression occurred (Ministry of Health 2005). The public health grading provided a formal assessment and public statement of the extent to which the treatment and distribution elements of a water supply system conformed with the DWSNZ, and whether adequate processes were in place to minimise the risk to public health (Ministry of Health 2003). Watercare's Statement of Corporate Intent (2005c) specified that the company had to maintain an 'Aa'<sup>18</sup> graded water supply system. The customer contracts specified additional quality and pressure requirements to be met at each BSP (e.g. maximum and minimum pressures), and the security of supply requirements for each bulk supply zone.

Watercare provided wastewater services to Auckland City, and the urban areas of Waitakere City, Manukau City, and Papakura district. No services were provided to Rodney District or North Shore City, which operated independent wastewater systems. Watercare's bulk sewer network received wastewater at designated locations from the LNO reticulation networks, and transported it to the Mangere Wastewater Treatment Plant for treatment and disposal by discharge to the Manukau Harbour.

Levels of service for wastewater were specified by resource consents and Watercare's contracts with the LNOs. Watercare's customer contracts specified the quantity of wastewater that Watercare's network had to be capable of receiving from the LNO networks at each collection point. The quality of treated wastewater discharges from the Mangere Wastewater Treatment Plant was governed by conditions specified in resource consents granted to the company under the Resource Management Act. Compliance with resource consent conditions was monitored by the Auckland Regional Council. The wastewater system also received discharges of liquid trade wastes from businesses in the region. Watercare was the regulating agency for those discharges, and as such had the power to impose conditions as might be warranted, and was responsible for monitoring compliance

---

<sup>18</sup> The 'A' refers to the water treatment grade; the 'a' refers to the water distribution grade. An 'Aa' grade was the highest achievable under the MoH's drinking water grading system.

with those conditions.

Watercare had no formal responsibility for the regional management of stormwater, but did receive stormwater flows into the bulk wastewater system. This occurred from two sources: from the wastewater collection system in Auckland City, parts of which were historically constructed as combined sewers (i.e. receiving both wastewater and stormwater), and from groundwater infiltration, illegal connections, and faulty or damaged pipes and manholes. Controlled overflows to streams and rivers around the region occurred from Watercare's wastewater network during peak storm events. Until 2001 these discharges were authorised under existing use rights issued under the Water and Soil Conservation Act 1967. After 2001, when those existing rights expired, Watercare worked with the LNOs and the Auckland Regional Council to progress the consenting of those discharges under the Resource Management Act.

**Table I.1.** Summary of Watercare's water and wastewater infrastructure, and budgets

Infrastructure	Statistics	Rep. Value
Water sources: 10 dams 1 groundwater source (Onehunga) 1 river source (Waikato River)	Total storage volume: 95 million m <sup>3</sup> 1-in-200-yr drought yield: 404,000 m <sup>3</sup> /day	\$310m
Raw water mains, tunnels, and aqueducts	79 km	
Water treatment plants (6)* * Hays Creek decommissioned in 2005	Total treatment capacity: 579,000 m <sup>3</sup> /day Ardmore: 333,000 m <sup>3</sup> /day	\$297m
Treated water mains	444 km (size range: 300 – 1900mm diameter)	\$815m
Service reservoirs	61	\$288m
Pump stations	Water: 30 Wastewater: 51	\$24m \$93m
Sewers	299 km (size range: 450mm – 2700mm diameter)	\$692m
Wastewater treatment plants	Mangere Wastewater Treatment Plant Treatment capacity: 360,000 m <sup>3</sup> /day	\$541m
Forecast Expenditure: Water (2009 – 2028, in 2007 dollars)	Capital Operational	\$1,164m \$429m
Forecast Expenditure: Wastewater (2009 – 2028, in 2007 dollars)	Capital Operational	\$1,566m \$833m

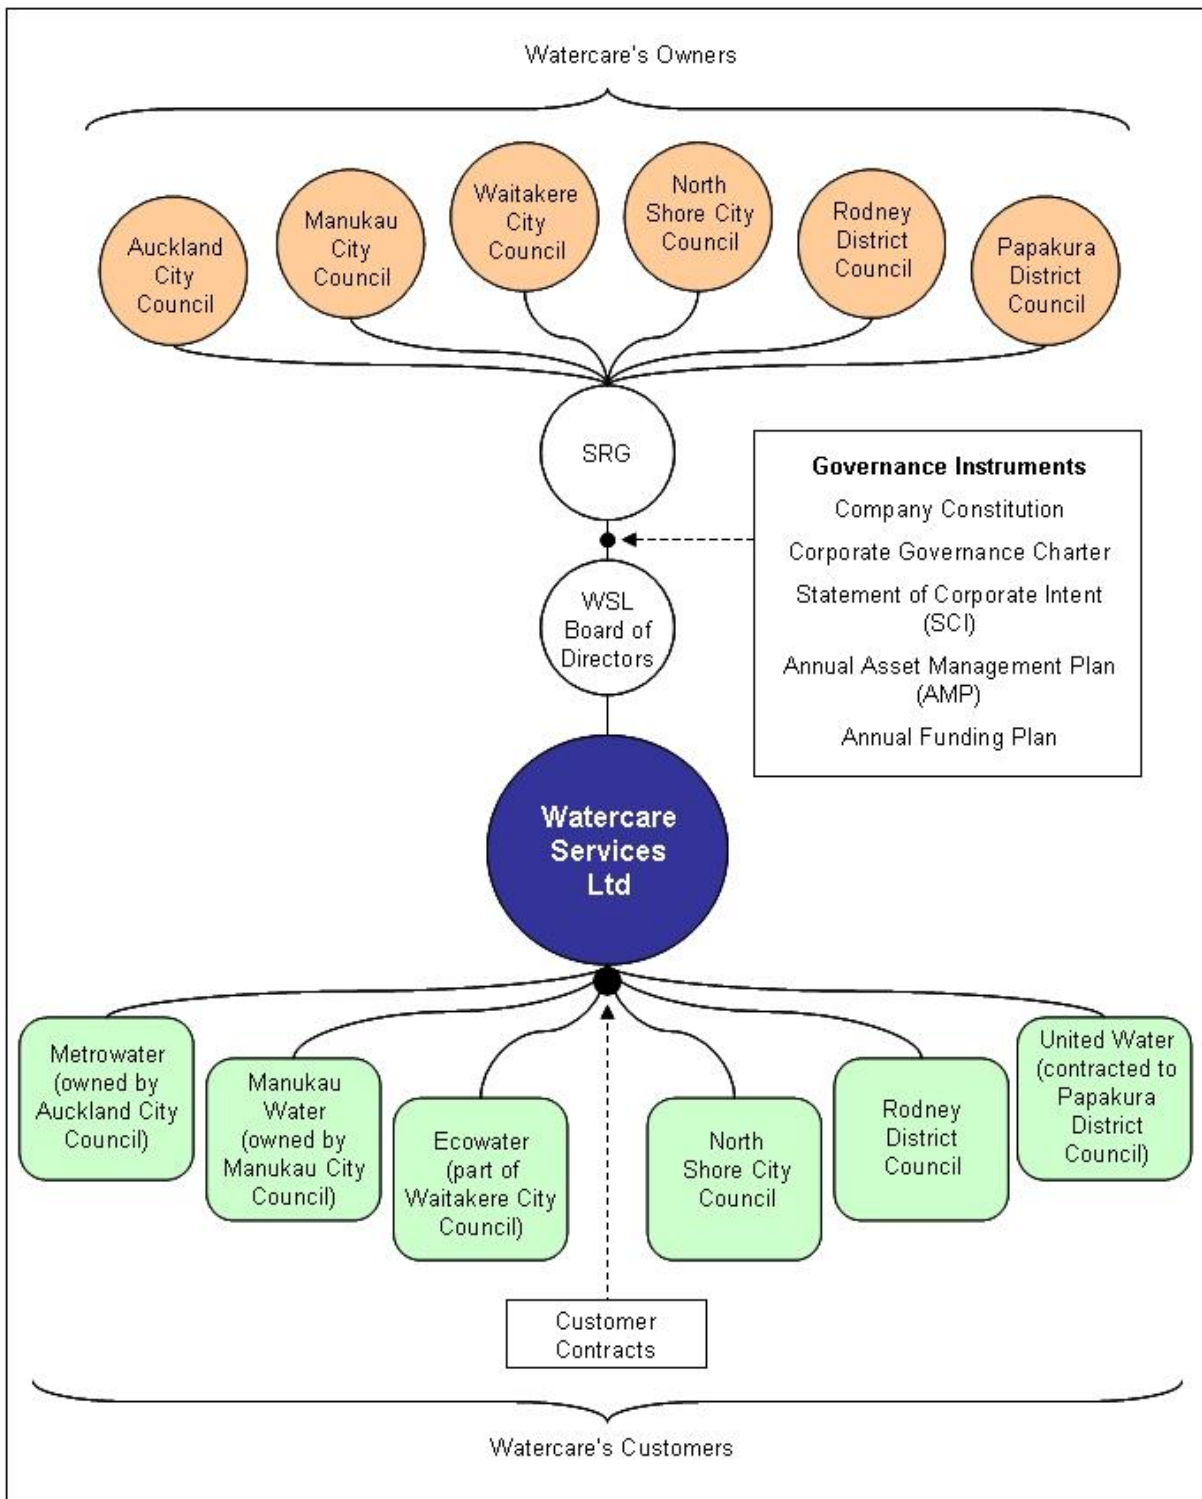
Sources for Table I.1: (Watercare Asset Management Plans, December 2005, December 2007)

## Corporate status and governance

Watercare operated under a mixed public-private governance arrangement. Historically the bulk water supply and wastewater assets of the company were owned and operated by the Auckland Regional Authority (predecessor of the Auckland Regional Council). Those assets were vested in Watercare in 1992 when the company was established as a Local Authority Trading Enterprise (LATE), owned by the Auckland Regional Services Trust (Watercare Services Ltd. 2005c, Appendix A). In 1998 the ownership of the company was vested in the city and district councils of Auckland, Manukau, North Shore, Papakura, Rodney, and Waitakere, and the company became a limited liability company registered under the Companies Act 1993 (Watercare Services Ltd. 2005c, Appendix A). The relative shareholdings were determined from the volume of water supplied to each city as a percentage of the regional total. This meant that Watercare's owners also owned or controlled Watercare's customers (i.e. the LNOs, with the exception of United Water in Papakura where the relationship was contractual). The industry structure is represented in Figure I.1.

Watercare's owners were represented by the Shareholder's Representatives Group (SRG), which performed high level governance functions on behalf of the owners, and was the formal vehicle through which Watercare's owners exercised their rights over the company (Watercare Services Ltd. 2005c, Appendix A). The SRG selected, appointed, and reviewed the performance of Watercare's Board, approved Watercare's annual Statement of Corporate Intent (SCI), considered the annual Asset Management Plans and Funding Plan on behalf of the owners, and approved any major acquisition or major transaction by the company on behalf of the owners (Watercare Services Ltd. 2005c, Appendix A). The Board of Directors was governed by Watercare's Constitution (2004b), which set out specific requirements and limitations, and a Corporate Governance Charter, which defined the duties and obligations of the Board in the areas of fiduciary duty, duty of care and diligence, and legal and statutory duties.

There was no independent price regulator for the Auckland Water Industry (e.g. as might be comparable with Ofwat in the UK). Retail prices for water services and charges for wastewater services were set by the individual LNOs and local government authorities. Watercare also set its own wholesale prices for water and wastewater services, subject to the conditions specified above. As was already noted above, the Ministry of Health regulated the quality of water supplies, while the Auckland Regional Council was responsible for regulating compliance with resource consents.



*Figure I.1. Former structure of the Auckland water industry (compiled from Watercare Services Ltd Annual Reports 2002, 2004, & 2007)*

## Mission and objectives

After 1 July 2003 when the Local Government Act 2002 (LGA 2002) came into force, Watercare was subject to dual legislative regimes: the “Council Organisation” provisions of the LGA 2002, and the Watercare specific provisions (Sections 707ZZZR and 707ZZZS) of the Local Government Act 1974 (LGA 1974), which were not repealed by the LGA 2002 (Watercare Services Ltd. 2005c, Appendix A). The LGA and Company Constitution (2004b) were the principal governing instruments of the company. The foundational provisions of these documents were as follows (*Local Government Act 1974; Local Government Act 2002; Watercare Services Ltd. 2004b*):

- Watercare must manage its business efficiently with a view to maintaining prices for water and wastewater services at the minimum levels consistent with the effective conduct of that business and the maintenance of the long-term integrity of its assets. This was Watercare’s primary legislative objective. It required the company to achieve the most efficient long term expenditure profile possible (referred to as operating at “least cost”) that maintained both the effective delivery of service (i.e. meeting the required service levels) and a tolerable level of risk over the long term. The trade-off between service, cost, and risk was therefore central to Watercare’s capital investment decision making.
- Watercare must not pay any dividend or distribute any surplus in any way, directly or indirectly, to its owners or any shareholder. This meant that Watercare could not make a profit, and prevented the company’s owners from using the company as a ‘cash cow’. The company could generate a surplus, but where this was not reinvested through the company’s Funding Plan, it had to be returned to the company’s customers (e.g. through rebates or price reductions).
- Watercare is limited to the performance of functions, and the conduct of business, in relation to waterworks, bulk water-supply, sewerage, and the treatment and disposal of sewage and tradewastes. This limited the scope of the company’s operations. In this regard, the company was granted the authority to exercise certain ancillary powers, including the provisions of the Auckland Regional Council Trade Waste Bylaw 1991 (for the regulation of trade wastes), and the powers of the Auckland Metropolitan Drainage Act 1960 (which provided powers for the construction, operation, and management of the sewer system).

- Watercare was prohibited from disposing of any assets which were necessary for the conduct of its business, from taking ownership of water services from a local authority or other local government organisations, and from entering into a contract of partnership with a local authority or other local government organisation relating to the operation of a water source without any time limitation on the matters that it could control. Similarly, Watercare's owners were prevented from selling their shares in the company (i.e. they had to remain owners). These provisions constituted and maintained the structure of the Auckland water industry (see Figure I.1).
- The local authority owners of Watercare were required to regularly undertake performance monitoring of the company to evaluate its contribution to the achievement (a) of their objectives for Watercare, (b) the desired results as set out in Watercare's Statement of Corporate Intent (SCI), and (c) the overall aims and outcomes of the various local authorities. Watercare's Constitution also required the company to prepare an SCI every year and within that document to specify, among other things, the objectives of the company and the performance targets and other measures by which the performance of the company may be judged in relation to its objectives. These provisions meant that, in addition to the principal objective specified in the LGA, Watercare was subject to such additional objectives as might be specified in the SCI, and that the company's performance against those objectives was to be monitored by Watercare's owners.
- The company Constitution stipulated that the principal objective of the Company shall be to operate as a successful business, which provides water and wastewater services that are economically viable, environmentally sound, socially responsible and responsive to customer needs. This statement established triple bottom line performance (i.e. Sustainability) as a central objective. The Statement of Corporate Intent specified a number of strategic performance objectives, collated under six Sustainability Policies: Environmental Care, Employees Health & Safety, Stakeholder Relationships, Customer Service Commitment, Asset Management, and Economic Performance. The company's performance was reported in the Annual Report using the same structure, although the link between the objectives in the SCI and the performance indicators used in the SCI was not entirely transparent (see Table I.2 which lists the SCI objectives for the period covered by this study, 2005 – 2008, against the corresponding performance objectives reported in the 2005 Annual Report).



**Table I.2.** *Watercare's corporate objectives as specified in the SCI and Annual Report*

Strategic Targets from the SCI (2005)	Objectives from Annual Report (2005)
<i>ENVIRONMENTAL CARE – To minimise the adverse impact of operations on the environment</i>	
<p>To ensure no successful prosecutions arise from breaches of Resource Management Act consents.</p> <p>To continue work on minimising the impact of biosolids and effluent on the environment.</p> <p>To co-ordinate the operations of the Water Advisory Group, in co-operation with the LNOs and ARC.</p> <p>To use energy efficiently and where appropriate recover energy from operational activities.</p>	<p>Promote cleaner production to industry and minimise waste</p> <p>Minimise the impact of biosolids and effluent</p> <p>Promote the preservation of species and protection of places of significant heritage value impacted by operations</p> <p>Use energy efficiently and, where practical, recover energy from operational activities</p> <p>Reduce and control odours, overflows, noise and other nuisances</p>
<i>EMPLOYEES, HEALTH &amp; SAFETY – To be an industry best workplace</i>	
<p>To demonstrate a commitment to best management practice in equal employment opportunities.</p> <p>To take all practical steps to provide employees with safe working conditions that do not detrimentally affect their health and safety and to have no successful prosecutions under the Occupational health and Safety legislation.</p>	<p>Provide staff with safe working conditions</p> <p>Maintain a working environment that promotes staff productivity and wellbeing</p> <p>Develop staff to their full potential</p>
<i>STAKEHOLDER RELATIONSHIPS – To be responsive to stakeholder requirements</i>	
<p>To work with the SRG and LNOs in a transparent and collaborative manner.</p> <p>To consult with the community and stakeholders on matters of relevance of them.</p> <p>To meet with the Maori Advisory Group and the Environmental Advisory Group at least quarterly.</p>	<p>Maintain open communication, educational initiatives and recreational opportunities</p> <p>Participate in relevant public policy and consultation with the community and stakeholders</p> <p>Comply with all statutory requirements</p>
<i>CUSTOMER SERVICE COMMITMENT – To provide high quality products and meet customer service level requirements</i>	
<p>To achieve full compliance with the Drinking Water Standards of New Zealand 2000 for potable reticulated water.</p> <p>To ensure water treatment plants operate to maintain the Ministry of Health (A) grading.</p> <p>To meet the 1 in 200 year drought security standards for the water supply system.</p>	<p>Continually improve service delivery to customers</p>
<i>ASSET MANAGEMENT – To manage and maintain the long-term integrity of assets</i>	
<p>To promote joint regional planning with all LNOs, including where appropriate co-ordination of capital investment decisions.</p>	<p>Ensure that capital projects have robust business cases and are delivered to plan</p> <p>Lead the development of an integrated vision for</p>

**Table I.2.** Watercare's corporate objectives as specified in the SCI and Annual Report

Strategic Targets from the SCI (2005)	Objectives from Annual Report (2005)
<p>To ensure Watercare's Asset Management Plans and processes are best appropriate New Zealand industry practice.</p>	<p>water demand, wastewater and stormwater management.</p> <p>Achieve continuous improvement in maintaining assets</p> <p>Achieve continuous improvement in sustainable business performance</p>
<p><i>ECONOMIC PERFORMANCE– To manage the business efficiently at minimum prices and to operate on a least-cost philosophy subject to fulfilling other environmental, social and legislative requirements.</i></p> <p><i>Note: In the Annual Report, this policy statement is shortened to: "To manage the business efficiently at minimum prices"</i></p>	
<p>To investigate and implement efficiency improvement opportunities within the water and wastewater industry in the Auckland Region, working closely with Watercare's existing customers.</p> <p>To maintain an "A" credit rating from the international rating agency Standard &amp; Poors and in so doing:</p> <ul style="list-style-type: none"> <li>Achieve a funds flow from operations to interest cover of 3.50 times or better.</li> <li>Renewing / replacing long-term debt facilities at least 6-months before their maturity.</li> <li>Achieve a net debt to total capitalisation ratio of less than or equal to 60% (excluding the impact of asset revaluations).</li> </ul> <p>To maintain the CPI-x pricing philosophy.</p> <p>Watercare will commission a targeted / focused cost efficiency review every 3 to 5 years as appropriate.</p>	<p>Ensure that financial strategies are consistent with achieving economic efficiency, intergenerational equity and an optimal cost of capital</p> <p>Ensure that the costs of providing services are consistent with leading practice</p>

## Company structure

Prior to 2004 the Watercare organisation was formally divided into two main business units, Water and Wastewater, with three ancillary units providing support services (design and project management, administration, and laboratory testing). The Water and Wastewater units were each organised along functional lines: planning (i.e. forward development planning), operations (i.e. asset operations and maintenance), and finance (i.e. accounting services). The company was restructured in 2004 to reflect the primary process functions of

the organisation. At the time of this study, the main business units were: Asset Management (responsible for asset planning and development, including project management), Operations (responsible for plant and network operations and maintenance), Finance (responsible for accounting and financial control, procurement, and IT services), and Business Services (responsible for administrative support, e.g. human resources, business planning). Support services were provided by the Office of the Chief Executive (oversight, audit, and communications), and the Office of the Chief Engineer (risk management, research and development, GIS). Despite the restructuring, however, there remained an informal separation between those who worked on the water supply side of the business, and those who worked on the wastewater side, largely because there were important differences between the two systems, which required the maintenance of distinctive technologies and knowledges.

## **Engineering asset planning and management**

Watercare existed by legislative mandate to provide water and wastewater services to the Auckland Region. The core business functions of the organisation were the development, operation, and maintenance of the infrastructure which delivered those services. Generally speaking, the operations function was concerned with the daily activities of operating the physical infrastructure systems. In modern treatment plants and water networks, most of these actions were automated, leaving human actors to perform inspection, oversight and corrective intervention (e.g. as performed by the central Control Room), and maintenance activities. Maintenance was concerned with the inspection, servicing, and rehabilitation of assets (planned maintenance), and with effecting repairs in the event of asset failures (unplanned maintenance). Forward maintenance planning sought to define a programme of inspection and maintenance activities which achieves a balance between the cost of the programme and the risk of asset failure. Together, the purpose of the operations and maintenance functions was to ensure the day-to-day delivery of services, the reliability of the company's existing infrastructure and operational capabilities, and recovery and restoration of those capabilities in the event of failure.

In contrast, the development function was concerned with the forward planning and construction of new assets to ensure that the company had the necessary infrastructure and operational capabilities to meet its service delivery requirements in the future. The planning process was conceptualised at two levels: strategic and master planning, and project

planning (see Figure I.2). In simple terms, strategic and master planning was concerned with the development of new capabilities at the system level. This type of planning sought to construct a comprehensive “birds-eye” view of the business environment, the trends within that environment, and the present and future performance of existing business capabilities, as a basis for identifying future capability shortfalls (usually understood as the ratio of the infrastructure capacity to the required demand and level of service). Strategic and master planning thus defined the need for (i.e. trigger) and timing of development projects. Project planning referred to the more detailed planning, investigation, and design work undertaken to develop a project from concept to construction (Figure I.3). At the end of the construction phase the new or refurbished asset was handed over to the Operations group to begin its useful life within the infrastructure system.

Inputs to strategic and master planning processes included the corporate objectives and specified current and future levels of service (defined in legislation and regulation, the statement of corporate intent, and customer contracts), demand forecasts (from the LNOs), existing asset condition information and performance histories, models of current and future infrastructure capacity (e.g. water network models, treatment plant process models), and a number of assumptions which simplified the analysis of the future operating environment. In addition to the renewal of existing infrastructure, the dominant drivers of Watercare’s investment programme, at a regional level, were the underlying trends of urban development, land-use change, and population growth across the region, regulatory changes affecting the company’s levels of service, and the idiosyncratic needs and demands of the LNOs, which operated in geographical areas differing significantly in their topography and urban forms, socio-economic diversity and growth rates, dominant industries, and political visions.

Formal strategic and master planning processes were not the only processes by which projects were identified. In principle, the trigger for a project could be anything, including master planning, condition assessments, other projects, operational concerns, failure events, customer requests, and so on. However, a trigger only became a formal project if, following an initial assessment, there was an identified business need or opportunity sufficient to warrant a project. The body of current and future projects identified through these processes constituted the Asset Management Plan (along with forward projections of operational, maintenance, and administrative expenditure). This body of projects was reviewed and prioritised on an annual basis. The Asset Management Plan (AMP) was the basis for Watercare’s annual Funding Plan and budgets. The annual preparation and publication of both plans was a legislative requirement.

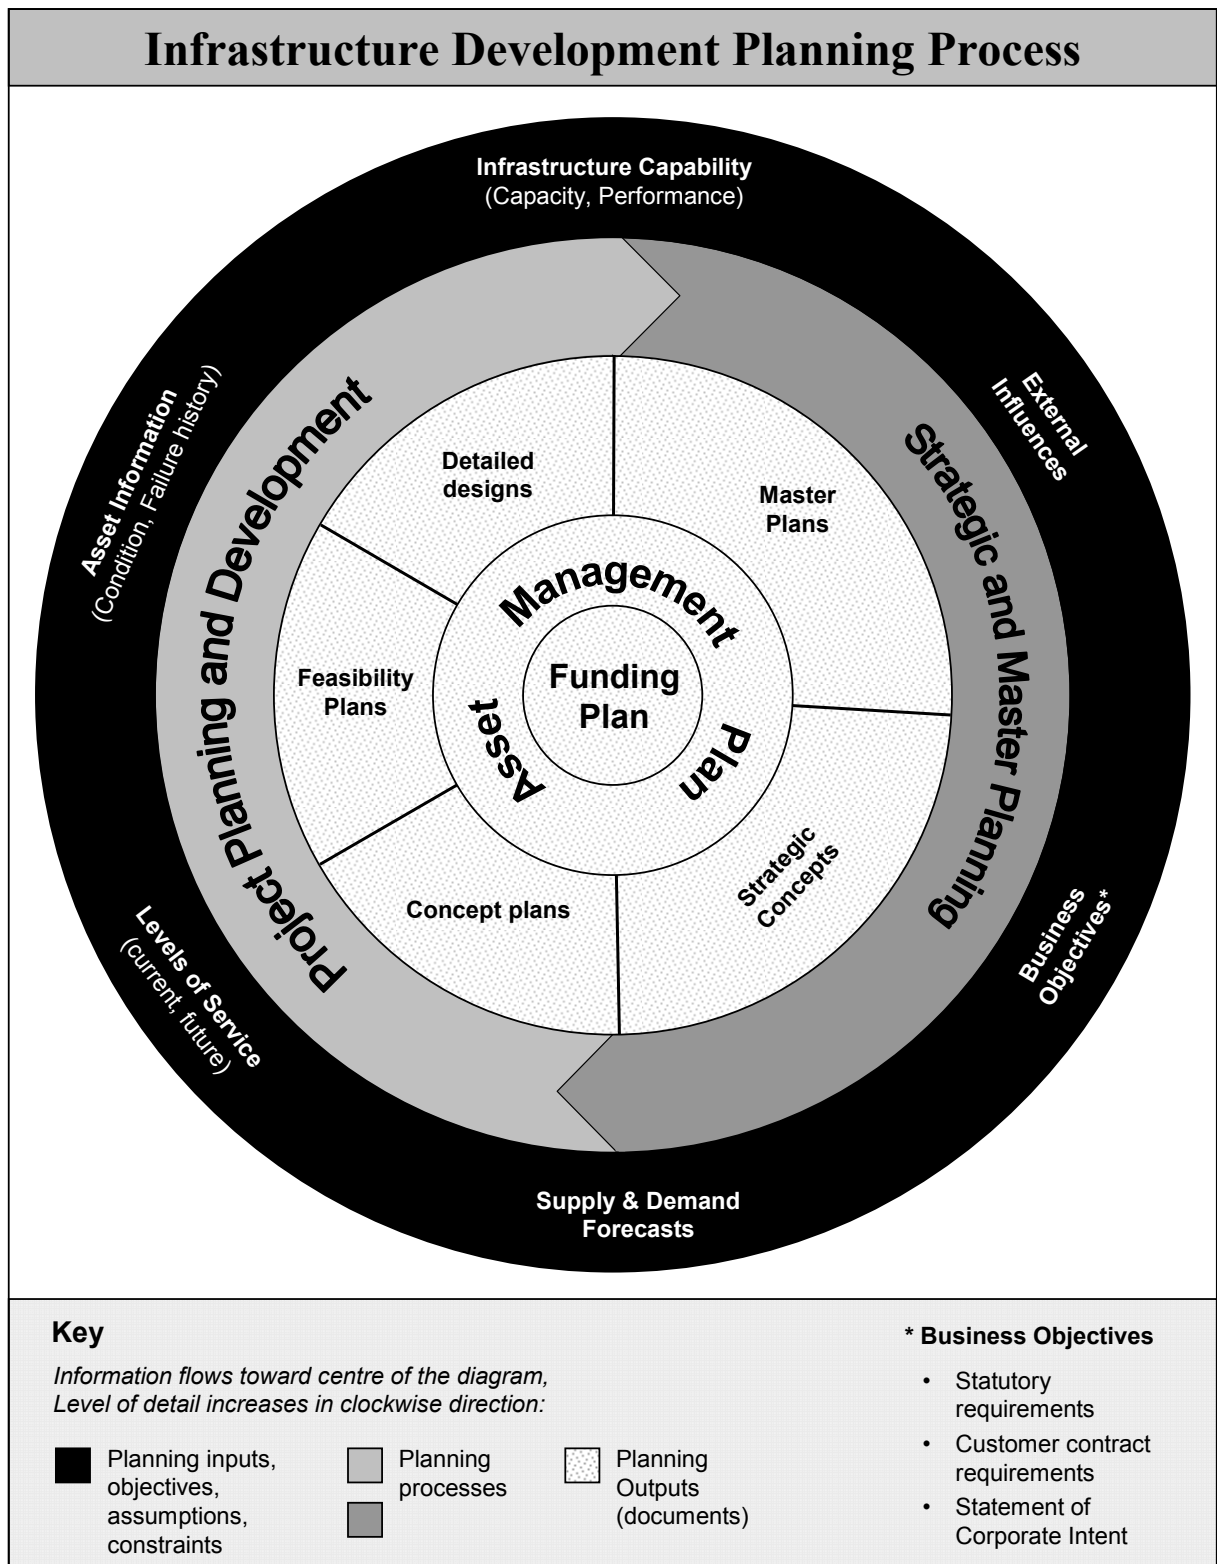
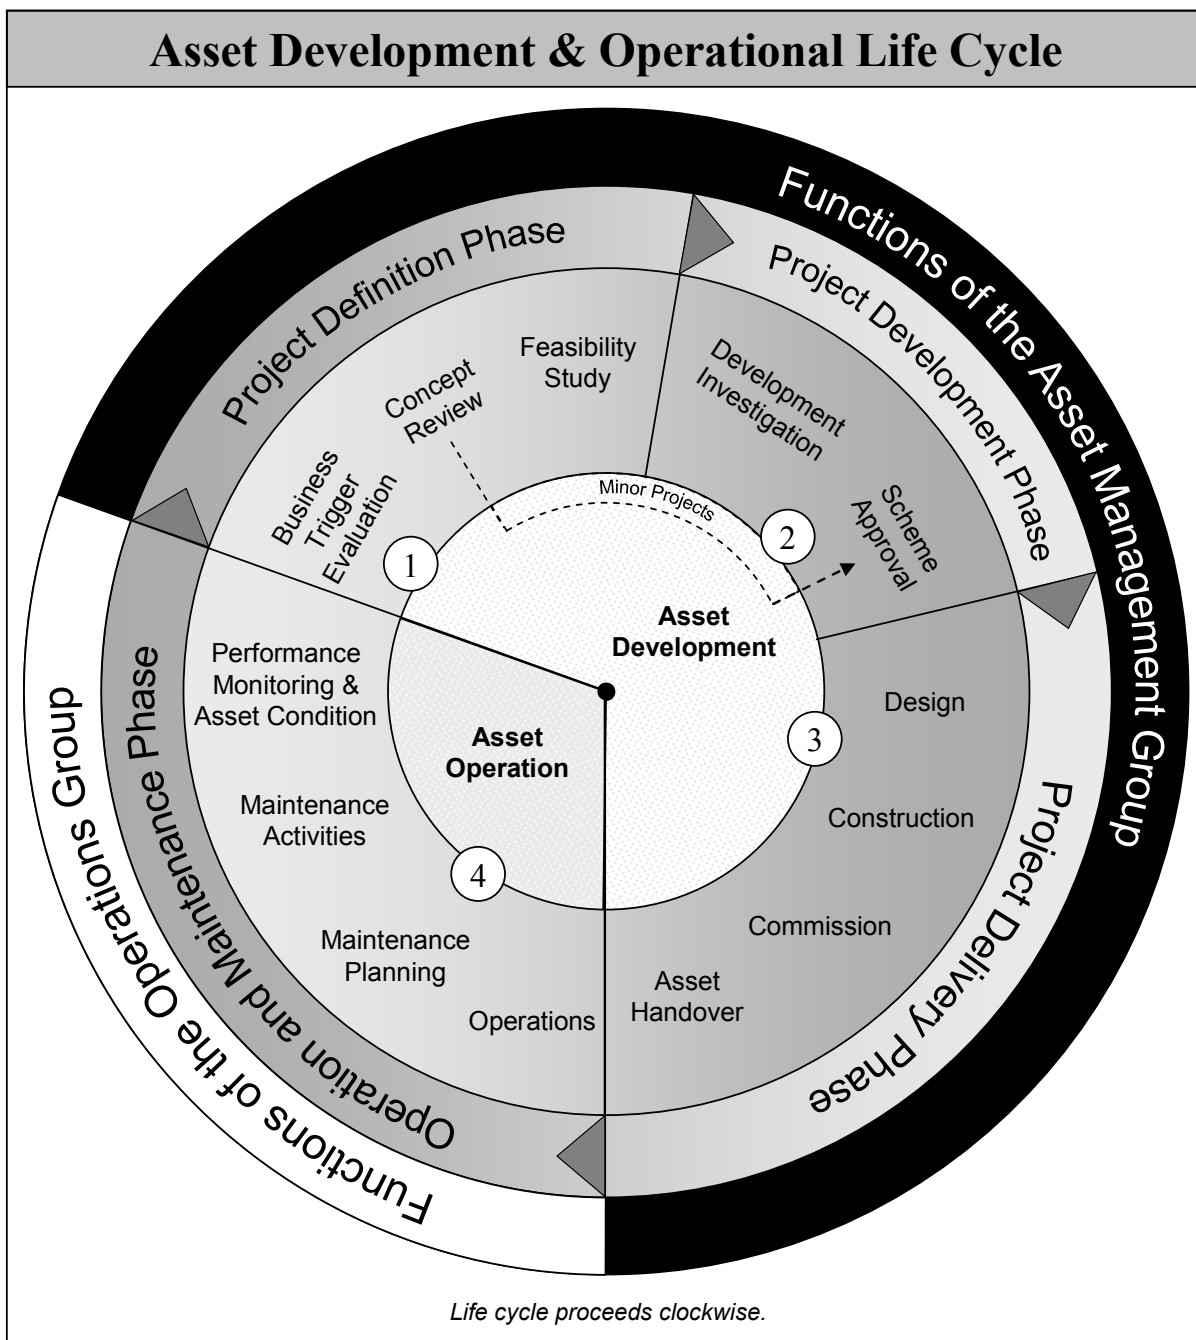


Figure I.2. Watercare's general infrastructure planning and development process



**Figure I.3.** The general development and operational life cycle of infrastructure assets. Numbers indicate key locations for use of Watercare’s risk assessment framework. Figure originally developed for draft of Watercare’s Project Development & Delivery Manual 2007

## Collaborative regional planning

Watercare did not carry out its asset planning in isolation. The company collaborated with the Local Network Operators on two levels. The first was necessitated by the fact that Watercare's infrastructure network was contiguous with the LNO networks. Thus, water quality and pressure in the LNO networks depended on the performance of Watercare's network, and, vice versa, changes in water consumption (demand) and operating requirements in the local area networks translated into specific level-of-service requirements at each of Watercare's bulk supply points (the formal boundaries between the Watercare and LNO networks). To manage this relationship, Watercare engaged with the LNOs both formally and informally throughout the annual asset management planning process.

The second level of engagement occurred over issues of regional importance or common interest amongst Watercare and the LNOs. Examples included the common framework for demand forecasting, regional demand management initiatives, or when Watercare undertook major projects potentially affecting several or all of the LNOs. Watercare hosted a monthly meeting of the Water Advisory Group, consisting of senior planners and managers from Watercare and each of the LNOs, to facilitate dialogue between the companies on common issues. More recently, Watercare also headed the Three Waters programme, a collaborative strategic planning initiative between Watercare, the LNOs, and the local councils, which sought to identify integrated long-term solutions for the management of water, wastewater, and stormwater across the region (Watercare Services Ltd. 2005d, 2008b).

Importantly, significant developments of the infrastructure networks occurred through these collaborative planning processes. Two projects are illustrative: the construction of the 38km Waikato River pipeline (completed in 2002) and the upgrade of the Mangere Wastewater Treatment Plant (completed 2003). The tapping of the Waikato River as a new source for Auckland's water supply was driven by the drought of 1994 which saw water restrictions enforced across the region. As a result of that drought, the local councils and network operators requested that the security of supply standard be raised from a 1-in-50-year drought to a 1-in-200-year drought (this being the international standard). The Waikato River was seen as a sustainable drought resistant source because the volume required for Auckland's supply would be less than 1% of the river flow at the abstraction point during a 1-in-100-year drought (i.e. the impact on the river would be negligible). The project also saw the introduction of new water treatment technology in the Waikato plant, considered to be the most sophisticated in Australasia. The project to upgrade the Manukau Wastewater Treatment Plant was launched in 1998 after five years of consultation and planning. The

project was driven not only by the need to expand the capacity of the plant to cater for Auckland's growth, but also by public and stakeholder aspirations for improved water quality in the Manukau harbour. The project saw the replacement of 500 hectares of oxidation ponds with an advanced, fully land-based, three stage treatment process, which reduced the wastewater treatment cycle from 21 days to 13 hours, allowed 13km of shoreline to be restored, and produces an effluent considered to be of "bathing water quality" (Watercare Services Ltd. 2002b). The Waikato and Manukau projects both produced significant step-wise changes in technology and the performance of the infrastructure networks, and both were the products of formal, top-down planning processes, involving extensive consultation with stakeholders and Auckland communities. The projects also shared another common feature. They were driven by regional concerns where significant performance objectives for the system at large were altered. Under business-as-usual decision making neither project would have been justifiable. But in each case a major performance objective was changed, necessitating a search for new solutions; for Waikato, it was the raising of the regional drought security standard, and for Manukau it was the de facto raising of acceptable standards for the quality of wastewater discharges to the Manukau harbour.

### **Capital decisions: the Service-Cost-Risk trade-off**

While the people and businesses of the Auckland Region were ultimately the final consumers of Watercare's services, as a bulk supplier the company's formal market consisted of only six customers, the LNOs, each of whom operated infrastructure networks that were contiguous with Watercare's, constituting a single uninterrupted network across the region. The legal relationship between Watercare and the LNOs, as that of wholesaler to retailer, is defined at certain physical points within the network by a mixture of regulatory and contractual specifications. But, Watercare's relationship with the LNOs, and indeed with the people and businesses of Auckland, was also clearly physical. What happened in Watercare's network affected what happened in the LNO networks, and ultimately also the end consumer. Although changes in the performance of Watercare's network were generally not of sufficient magnitude to be noticeable to the end consumer under normal operating conditions (i.e. minor daily fluctuations in pressure and quality), this context required the company to maintain an extraordinarily high degree of reliability. Operational failures could have very real and very dire consequences far beyond the impacts to the company's reputation or financial statements.

The importance of Watercare's services to the regional economy, and to the health and



wellbeing of the people of Auckland was reflected in the stringent regulatory quality standards that the company had to meet, and in the company's unique governance regime. Although constituted as a private, limited liability company, Watercare's shareholders were public sector organisations (the local councils of the six cities and districts served by Watercare), and the company's existence, ownership structure, and scope of operations were fixed by legislative mandate. Recognising that Watercare's services constituted essential public goods, that mandate expressly forbid the company from making a profit, but still required the business to operate as efficiently as possible "with a view to maintaining prices for water and wastewater services at the minimum levels consistent with the effective conduct of that business and the maintenance of the long-term integrity of its assets". This legislatively defined objective required the company to achieve the most efficient long term expenditure profile possible that maintained both the effective delivery of service (i.e. meeting the required service levels) and a tolerable level of risk over the long term. The trade-off between service, risk, and cost was therefore central to Watercare's capital investment decision making.

## The risk assessment framework

Watercare began to develop a formal risk management framework in 1998 with advisory input from consultants Broadleaf Capital International Australia Ltd and Woodward-Clyde (NZ) Ltd (Cooper et al. 1998). The framework and risk management process were based on the then current Australia/New Zealand Risk Management Standard, AS/NZS 4360:1995. Individual spreadsheet-based risk registers were initially maintained for each of the main business units of the company (Water, Wastewater, and Corporate), but these were amalgamated into a single corporate risk register (a Microsoft Access database) in 2001. A separate project risk register was developed in 2003 in the form of a spreadsheet template for use by project managers. The formal policy, scope, and objectives for risk management, as they stood at the time the new Corporate Risk Manager joined the company in early 2007, are shown in Figure I.4.

The Risk Management Framework manual (Watercare Services Ltd. 2003c, p3) stated that the formal responsibility for risk management lay with the Chief Executive, being "responsible the establishment, implementation, and maintenance of sound risk management practices in the organisation". That responsibility was delegated to the company's managers, being responsible for "ensuring that the risk management processes as

defined in the Risk Management Framework are implemented effectively in their areas of responsibility”, and to Watercare’s staff, being responsible “for identifying, analysing and managing risks in their areas of activity” (Watercare Services Ltd. 2003c, p 3). The Risk Management Steering Committee, made up of the various company general managers, was formally responsible for administering the company’s risk management policy and for oversight of the risk management function, although this responsibility was delegated to the Internal Auditor for day-to-day management.

Oversight of risk management was formally separated from internal audit in 2004 with the creation of a dedicated Corporate Risk Manager role reporting to the Chief Engineer. The period 2004 to 2006 was one of transition within both the CRM and Chief Engineer roles, with each role being held by three people over that time (a result of incumbents moving on to pursue opportunities elsewhere, and internal reshuffling). The third CRM, and subject of the study reported in this thesis, was appointed in early 2007.

Policy, scope, objectives for Risk Management at Watercare (c. 2007)	
<b>Policy:</b> ...to manage risk in a prudent manner to enable business objectives to be consistently met.	
<b>Scope:</b> ...to ensure that: <ul style="list-style-type: none"> <li>▪ Risk is understood and identified;</li> <li>▪ Hazards and practices that could cause financial loss, disruption to business goals, injury to people or damage to the environment are identified and controlled as far as practicable;</li> <li>▪ Insurance or other financial arrangements are established to protect the business interests should a loss damaging to the finances of the business occur.</li> </ul>	<b>Objectives:</b> ...to provide: <ul style="list-style-type: none"> <li>▪ Protection and continuity of core business activities;</li> <li>▪ Fulfilment of legal obligations;</li> <li>▪ Safeguards for public and employee health;</li> <li>▪ Environmental protection;</li> <li>▪ Operation and protection of assets at lowest cost;</li> <li>▪ Contingency planning for foreseeable emergency situations;</li> <li>▪ Protection of the balance sheet.</li> </ul>

*Figure I.4. The formal policy, scope, and objectives for Risk Management at Watercare (source: WSL Risk Management Framework manual 2003, p 3).*

## Scoring and classification

Two revisions were made to the risk scoring and classification framework between 1998 and 2007. A simple qualitative risk matrix was originally used for categorising risk consequences, but this was eventually found to have inadequate definition for differentiating and prioritising the large body of risks that had been identified. The first revision, in 2003, saw this evolve into a semi-quantitative five-by-six matrix consisting of five consequence categories (Reputation, Finance, Environment and Public Health, Health and Safety, and Asset Management) and six severity categories (from Negligible to Catastrophic). The various levels

of severity were defined by qualitative descriptors in each consequence category and assigned a score between 0 and 100. Similarly, the likelihood scale was defined by qualitative descriptors, from remotely possible (less than one occurrence in 100 years) to almost certain (more than one occurrence per year), with each gradation of the scale also assigned a score between 0 and 100. In risk assessments, a root-mean-squared (RMS) averaging method was employed to reduce the five-dimensional consequence scale to a single score. The move to a semi-quantitative risk scoring method allowed that all risks could be prioritised on the basis of their numerical risk score (between 0 and 10,000), and represented diagrammatically on a Likelihood vs Consequences chart.

Except for the RMS averaging method, Watercare's risk classification and scoring framework was consistent with the methods described in AS/NZS 4360. The two-dimensional risk space defined by the Consequences vs Likelihood chart was divided into five risk classes, with the boundaries between the classes defined as lines of constant risk score. The location of those boundaries was established according to how a group of experienced staff perceived relative asset risk priorities (Clement 2007e). They grouped various risks using qualitative descriptions of the different risk classes as guides. The risk scores that separated those groupings were then chosen to define the risk classification boundaries. Formally, the company's policy, as stipulated in the 2003 risk framework manual, was that Class 1 risks were considered "Acceptable" while Class 5 risks were generally considered "Unacceptable", with risks falling in classes 2 to 4 being considered tolerable as long as they were managed to be As Low As Reasonably Practicable (ALARP). It may be noted, however, that while this is generally consistent with the definition of ALARP given by the The Health and Safety Executive of the UK (HSE 2008), Watercare's Risk Management manual did not provide any guidance on how to make an ALARP judgement. The second revision of the classification and scoring framework in 2006 saw a slight change to the definition of those risk classes, which was motivated by a desire to reduce and rationalise the number of priority risks (Classes 4 and 5) reported to the Board. The remaining priority risks were then associated with a limited number of critical assets, which were called the Enterprise Risks and from then on formed the basis for risk management reporting to the Board (Clement 2007e).

## **Use of the risk framework**

By early 2007 the risk management framework had been in use within Watercare for over eight years. At that time, the risk scoring and classification framework was employed in five primary ways. The numbers, below, correspond to the numbers on Figure I.3, indicating key

locations for use of the risk assessment framework within the asset management lifecycle.

- i) Risk reviews: Risks in the corporate risk register were reviewed on a regular (6-monthly or annual) basis depending on the risk class. This included identifying and assessing new risks, and updating existing risk assessments. Where new risks were identified but could not be addressed via operational or maintenance controls, then they became triggers for capital projects.
- ii) Capital decision making: The risk assessment framework was used to evaluate the business need for capital investment. This occurred formally in written capital expenditure requests (CapExs), which required approval by the Chief Executive (and sometimes the Board, depending on the magnitude of the expenditure). Each CapEx had to present a range of options for addressing the identified business need, each of which was evaluated in terms of cost and risk. The cost of the option was the net present value (NPV) of the capital expenditures plus any operational expenditure changes attributable to the option, typically assessed over 20 years. The risk was calculated as the net risk reduction between the status quo (or do minimum) and the option. As a general rule, the option with the best risk reduction to cost ratio was the preferred option (i.e. the option that reduced the risk to an acceptable level or which achieved the most economical level of risk reduction). The risk reduction was represented diagrammatically on the above likelihood versus consequences charts.
- iii) Project management: A separate project risk register was used by project managers to identify, record, and manage risks on all projects over NZ\$2 million capital value, and on projects of lesser value where warranted by the degree of risk involved. A simpler, standardised template of common project-related risks was also utilised as a management tool on smaller projects. Both templates employed the risk classification and scoring framework described above.
- iv) Maintenance programme optimisation: In 2005 the company initiated a Reliability Centred Maintenance (RCM) regime to improve the programming of planned maintenance activities. This involved the computer modelling of the reliability of plant, equipment, and processes to predict an optimum frequency for planned maintenance interventions (i.e. routine inspections and refurbishments). Standard industry models of failure frequency (Weibull curves) were modified using recorded equipment-specific failure and maintenance histories, and used to predict process reliability. Planned maintenance interventions were then programmed to optimise

the risk (of failure) versus cost (of maintenance) trade-off. The company's risk classification and scoring framework was built into the RCM models for this purpose.

- v) Reporting to the Board: The above classification framework and charts were used to represent risk magnitudes with reporting to the board. In addition to key operational and project risks, a small number of Enterprise Risks were regularly reported. These Enterprise Risks were collations of priority risks associated with critical assets within the infrastructure network.

## **Significance and quality of the Risk Management function**

Watercare's public discourse consistently promoted risk management as an integral part of the company's operations, and risk as a main factor in corporate decision making. Prior to 2006, Watercare's annual reports highlighted risk management as a key governance mechanism, as the basis for Board and management decision-making, and as fundamental to the achievement of the company's core legislative purpose, while the revision of the risk assessment framework in 2003 was considered significant enough for the Chairman of the Board to comment on it in the 2004 Annual Report (see excerpts in Figure I.5). Prior to 2003 risk-based decision-making was also listed as a corporate objective and reported on under the Statement of Service Performance at the back of the company's annual reports. The 2000 Annual Report, for example, reported that the objective that all capital projects should be justified on the basis of "cost, risk, business need, and impact on earnings" had been met (Watercare Services Ltd. 2000, p 78). From 2003 onwards, however, this practice was gradually phased out. In the 2004 and 2005 annual reports the reported risk management objective was only to ensure that all risk analyses on projects over \$NZ10 million in capital value were independently reviewed (Watercare Services Ltd. 2004a, p 101; 2005a, p 92). With the exception of statements from independent reviewers, the 2006 and 2007 annual reports did not comment or report on the company's risk management function. The company's Asset Management Plan's, however, continued to describe how risk management featured in planning and decision-making. The 2006 AMP, for example, stated that risk management was an integral part of managing the life cycle of major infrastructure assets and drew attention to the identification of all significant risks to the business, the assessment of those risks in terms of the aforementioned multi-dimensional consequences framework, the risk-based prioritisation of business projects, and the active management of risks through replacement, rehabilitation, monitoring, and contingency planning (Watercare Services Ltd. 2006, p 33).

<b>Public comments made by Watercare about the quality of the company's risk management</b>	
<p>“Watercare assesses significant corporate decisions in light of the degree of risk to which they expose the company and the assets and services for which the company is responsible. In other words, the company actively manages its risks.”</p> <p>(WSL Annual Report 2003, p 31)</p>	<p>“Watercare modified its risk management framework during the year to give better balance to environmental, social, and economic factors in the decision-making process. Several initiatives – including the cell lysis technology reported on page 32 and the Hobson Bay tunnel proposal reported on page 13 – were progressed partly because of their strong environmental and social benefits. The framework also helps the Board assess the contribution major projects make towards sustainability.”</p> <p>(WSL Annual Report 2004, p 5)</p>
<p>“Watercare’s governance structure is supported by a risk management framework and a decision making process that ensure efficient results-based management, a transparent flow of information to shareholders, effective risk mitigation and leading-practice risk assurance.”</p> <p>(WSL Annual Report 2004, p 15)</p>	<p>“Risk management-based planning ensures Aucklanders can depend on Watercare to deliver its services. It also assures Watercare’s owners that the business is managed prudently. Potential risks are assessed for their probability and multidimensional consequences... The system provides the company and its board with a semiquantitative risk-priority rating that supports informed decision-making.”</p> <p>(WSL Annual Report 2005, p 19)</p>
<p>“Watercare’s risk management framework embraces the fundamental role the company plays in supporting the region’s economy and the lifestyle of its residents... As Watercare’s primary responsibility is to maintain essential water and wastewater services to the people of Auckland, it is appropriate that risk management forms the framework for the company’s decision-making process.”</p> <p>(WSL Annual Report 2004, p 18)</p>	

**Figure I.5.** Public commentary on risk management in WSL Annual Reports 2003 - 2005  
(sources: Watercare Services Ltd. 2003a, 2004a; 2005a, page numbers as indicated)

The excellence of the company’s risk management framework was also consistently promoted. In addition to providing verification that the company’s risk management framework complied with AS/NZS 4360, various external reviewers commented, in their letters of opinion (published in the company’s annual reports), to the effect that risk management at Watercare could be considered leading practice. It was reported in the 2002 annual report, for example, that the review by Broadleaf Capital International had confirmed that “Watercare’s risk management procedures and practices are consistent with international best practice” (Watercare Services Ltd. 2002a, p 13). Lane and Associates expressed a similar opinion in the 2005 annual report:

The framework provides a robust risk identification and management system that should provide Watercare with a high level of confidence that its significant risks are being identified... The risk-management process is well understood and is generally well communicated. The framework’s

application, and the organisation's dedication to further advance risk-management practices throughout its activities can be considered leading practice. (*excerpt from Watercare Services Ltd. 2005 Annual Report, Risk Management Review, p25*)

## Key performance statistics for the period 2000 - 2008

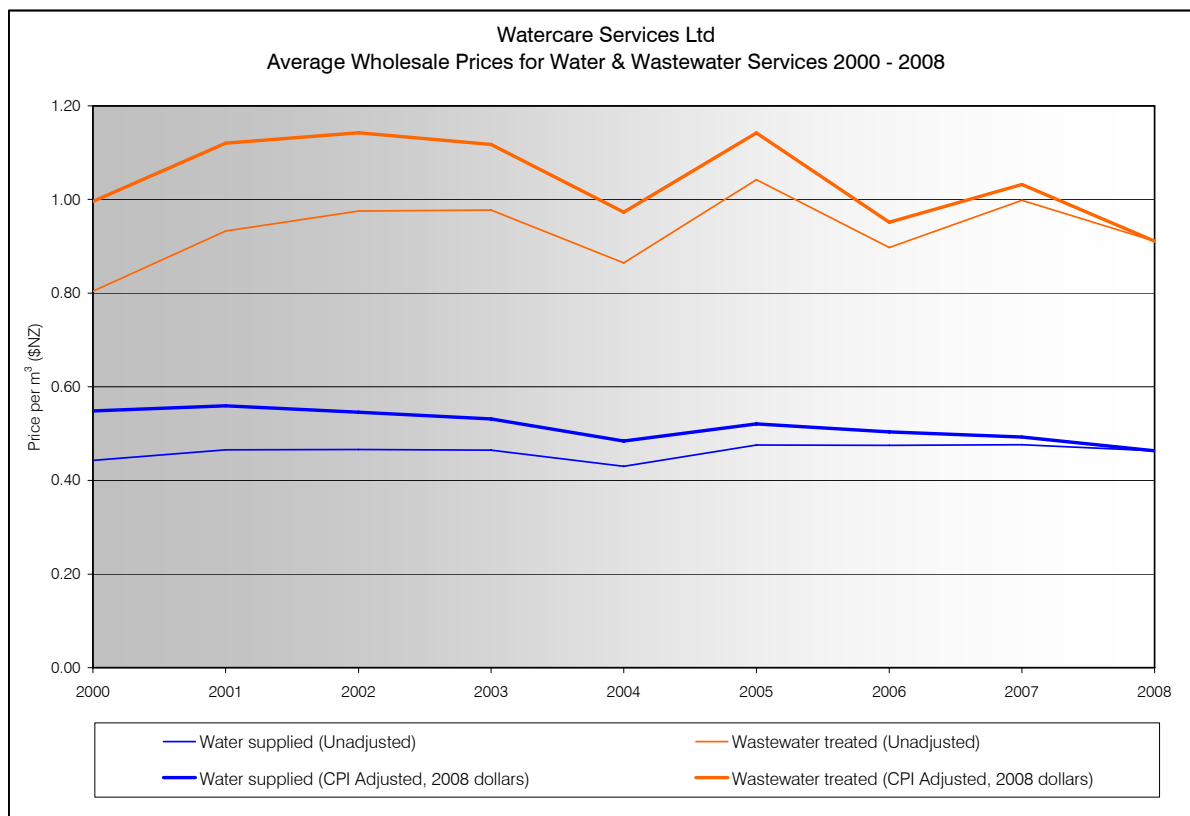
Watercare was held accountable against a wide range of performance indicators, most of which were reported or at least summarised in the company's Asset Management Plans and Annual Reports. For reasons of practicality it is clearly not possible to summarise every facet of Watercare's performance here, but the general nature of that performance may be discerned by examining a representative cross-section of indicators. A number of enterprise performance indicators representing key dimensions of the company's activities are summarised in the following sections for the period 2000 – 2008 (see Table I.3 below). The indicators presented include both direct measures of performance (e.g. meeting service delivery standards, customer satisfaction, etc.) and measures of process and system quality (e.g. the MoH grading, the WSAA benchmark scores, the Standard & Poors credit rating). All of the indicators presented in the following section were drawn from Watercare's public Annual Reports and Asset Management Plans. The performance trends are summarised in the last section below.

*Table I.3. Dimensions of Watercare's performance and corresponding key performance indicators*

<b>Value-For-Money</b>	Wholesale prices for water and wastewater services.
<b>Water service delivery</b>	Non-compliance against the NZ Drinking Water Standards. Maintenance of the drought security standard
<b>Wastewater service delivery</b>	Non-compliance against resource consents for the discharge of effluent from the Mangere Wastewater Treatment Plant. Percentage of overflows from the wastewater collection network not attributable to wet weather.
<b>Quality of business processes</b>	Variance against operational and capital expenditure budgets WSAA asset management process benchmark score Maintenance of the MoH 'Aa' grading S&P credit rating
<b>Health &amp; Safety</b>	Lost-time injury frequency rate

## Value-For-Money

There was no standardised measure of Value-For-Money in the Auckland Water Industry. A proxy measure is the prices charged by Watercare for its services (presented in Figure I.6) considered against the quality and reliability of those services (see other indicators).



**Figure I.6.** Watercare's wholesale prices for water and wastewater services 2000 – 2008.

Source: calculated from annual turnover versus volumes supplied/treated quoted in Watercare Annual Reports 2000 – 2008; cf. also Figure 63, p 144 of the 2008 AR. CPI adjustments calculated from CPI figures obtained from Statistics NZ ([www.stats.govt.nz](http://www.stats.govt.nz))

## Non-compliance with regulatory standards for service delivery

The primary performance standards for water service delivery are the NZ Drinking Water Standards and the various conditions of Watercare's contracts with its customers. The degree of non-compliance with these conditions is therefore a key measure of water service delivery performance. Similarly, the primary performance standards for wastewater treatment are the resource consent conditions imposed on the Mangere Wastewater Treatment Plant and some parts of the wastewater collection system. The degree of compliance with those conditions is therefore a key measure of wastewater service delivery performance (Watercare's Corporate Risk Manager adopted the degree of compliance with



statutory obligations as a key indicator of the achievement of Watercare's corporate objectives for his risk framework). Table I.4, below, summarises Watercare's performance with respect to compliance with the NZ Drinking Water Standards and compliance with resource consents for discharges from the Mangere Wastewater Treatment Plant.

**Table I.4. Watercare performance: Non-compliance with regulatory standards 2000-2008**

Year	Non-compliance with NZDWS	Non-compliance with resource consents for Mangere WWTP discharge
1999/00	None [p 81]	Eight breaches of ammonia, nitrogen, faecal coliforms and enterococci (due to limitations with operation of the old plant) [p 81-82]
2000/01	None [p 79]	Enterococci limit exceeded Dec/Jan BOD limit exceeded Jan, Feb, March. [p 25, 84]
2001/02	None [p 49, 110]	None [p 106]
2002/03	Total of one breach of Turbidity standard at Huia [p 131]	3 breaches of non-filterable residues consent in Dec/Jan/Feb due to salt crystallisation from infiltration of seawater to the intertidal storage basin as a result of efforts to control midges. [p 127]
2003/04	Total of three breaches of Turbidity standard at Huia and Waitakere plants. Breaches quickly resolved with no measurable health effects, and no effects on Plant gradings. [p 104 + Fig. 50 in Additional Data]	Suspended solids limit exceeded in four months, and minimum UV dose not achieved in three months due to carry over from reactor clarifiers (rectified by construction of thickening plant), and dredging of intertidal storage basin. [Fig. 42 in Additional data]  24 technical breaches (minimum dissolved oxygen, non-filterable residue, minimum UV dose, maximum BOD, minimum pH, total petroleum hydrocarbons). [p 47 + 102]
2004/05	None [p 95 + Fig. 42, p 121]	20 technical breaches on minimum pH, non-filterable residue, and levels for petroleum hydrocarbons, nitrogen, and ammonia levels. No detectable effects. [p 93 + Fig. 41, p 120]
2005/06	None [p 91 + Fig 42, p 117]	Some minimal exceedences of ammonia (Feb and June) and nitrogen (Jan and Feb) levels and minimum UV dose not achieved in June. [Fig. 41, p 116]
2006/07	None [p 95, 129]	Total 7 days minimum UV dose not achieved due to stormwater flows. [p 128]
2007/08	None [p 138]	Total 10 days minimum UV dose not achieved due to stormwater flows. Total 34 days ammonia levels exceed in final effluent due to wet weather flows. [p 137]

Sources for Table I.4: The source for each year is Watercare's Annual Report for that year [page numbers in square brackets].

## Asset management processes

The degree of variance against the annual operational and capital expenditure budgets can be taken as a measure of operational performance (i.e. significant and repeated variance against these budgets would indicate problems with core business planning and management processes). The CRM adopted the achievement of the budget and the AMP as key indicators of the achievement of the corporate objectives for his risk framework. Table I.5 summarises Watercare's OPEX and CAPEX variances for those years where they were reported in the company's Annual Reports.

*Table I.5. Watercare performance: Variance against Operational and Capital Expenditure Budgets*

Year	Percentage variance against operational budget	Percentage of forecast capital expenditure completed
1999/00	NR	NR
2000/01	NR	NR
2001/02	NR	95% [04AR Fig. 9, p 20]
2002/03	+5.4% [p 91]	80% [04AR Fig. 10, p 20]
2003/04	-2.0% [p 60]	>95% [p 53]
2004/05	+1.5% [p 58]	81% [p 51]
2005/06	+1.5% [p 50]	81% [p 43]
2006/07	-3.7% [p 58]	90% [p 49]
2007/08	-2.2% [p 56]	84% [p 47]

*Sources for Table I.5:* Unless otherwise indicated the source for each year is Watercare's Annual Report for that year [page numbers in square brackets].  
NR = Not Reported.

Since 2004 Watercare has annually self-assessed its asset management processes against the Water Services Association of Australia's (WSAA) asset management benchmarking programme. The company's score against this benchmark may be interpreted as indicator of the quality of the company's asset management capabilities. In 2008 the WSAA benchmarking project was combined with the similar Aquamark project administered by the International Water Association. In 2004 Watercare's scores across the seven evaluation categories ranged from 55 to 98, while in 2008 the score range was 75 to 95, with all but one score over 80, indicating an overall improvement (Watercare Services Ltd. 2008a, p 55).

## Other

Additional measures of Watercare's performance are presented in Table I.6.

*Table I.6. Watercare performance: Miscellaneous 2000-2008*

Year	1-in-200-year drought security standard	Ministry of Health 'Aa' grade for water supply	S&P credit rating	Lost Time Injury rate* <sup>2</sup> (LTI per million hours worked)	% overflows from WW network not due to wet weather
1999/00	Yes [p 81]	Yes [p 81]	A+/A-1 [p 81]	12.1	0.5% * <sup>3</sup>
2000/01	No [p 79]	Yes [p 79]	A+/A-1 [p 49]	12.1	1.1% * <sup>3</sup>
2001/02	No [p 110]	Yes [p 110]	A+/A-1 [p 62]	10.9	1.2% [p 38]
2002/03	Yes [p 131]	Yes [p 131]	A+/A-1 [p 77]	9.4	1.2% [p 51]
2003/04	Yes [p 104]	Yes [p 104]	A+/A-1 [p 101]	7.2	0.9% [p 38]
2004/05	Yes [p 95]	Yes [p 95]	A+/A-1 [p 92]	5.8	1.0% [p 113]
2005/06	Yes [p 91]	Yes [p 91]	A/A-1 [p 89]	0.0	1.8% [p 109]
2006/07	Yes [p 95]	Yes [p 95]	A/A-1 [p 93]	1.3	1.5% [p 121]
2007/08	Yes [p 101]* <sup>1</sup>	Yes [p 101]	AA/A-1+ [p 97]	5.0	1.8% [p 129]

Sources for Table I.6: Unless otherwise indicated below the source for each year is Watercare's Annual Report for that year [page numbers in square brackets].

\*<sup>1</sup> Changed to the international standard of 1-in-100 years in 2008.

\*<sup>2</sup> Reference for all LTI figures is the 2007/08 AR, Fig. 38, p 131.

\*<sup>3</sup> Also sourced from 2001/02 AR, p 38.

Notes to Table I.6:

- *Drought security standard:* The drought security standard is a statement of the region's tolerance for drought risk. Compliance with this standard is a measure of the risk of supply disruption due to extreme drought conditions.
- *MoH Grading:* The Ministry of Health grading of public water supply systems is a risk-based assessment of the quality of those systems. It takes into account the sophistication of treatment barriers in place, the age, and condition of infrastructure, the quality of inspection and monitoring systems and procedures, the quality of management, assurance, and response procedures, the experience of supervision personnel, and the frequency and severity of performance transgressions. The grading given to Watercare's water treatment plants and supply network is therefore a measure of the quality of Watercare's infrastructure and management processes.
- *Credit rating:* The Standard & Poors credit rating is an internationally recognised measure of the credit worthiness of an organisation. The rating given to Watercare is therefore a measure of quality of the company's financial management practices.
- *Lost-time injury frequency:* The LTI frequency is a well recognised indicator of the efficacy of an organisation's health and safety procedures.
- *Wastewater overflows not due to wet weather:* Due to the way Auckland's sewer network was historically constructed the wastewater system receives considerable wet weather flows. The percentage of overflows from the network not attributable to wet weather flows is therefore a proxy measure of reliability and capacity of the system.

## Summary of Watercare's performance 2000-2008

Table I.7 summarises the trends in Watercare's performance over the period 2000-2008, which may be inferred from the indicators presented above. On the basis of Table I.7, I think it is reasonable to assert that Watercare consistently achieved a high level of performance over that period. It is also worth noting that between 2000-2008 Watercare delivered a capital programme totalling nearly NZ\$880 million, including the completion of the Mangere Wastewater Treatment Plant upgrade (\$450m), the Waikato Water Treatment Plant and pipeline (\$155m), the upgrades of the company's water treatment plants to meet the 2005 NZ Drinking Water Standards (\$50m), and commencing construction of Project Hobson in 2007 (\$111m budgeted in 2006 AMP). The upgrade of the Mangere WWTP (commenced in 1998) and the Waikato pipeline (commenced in 1999) together represented the largest capital expansion in Watercare's history.

Of course, the company's record was not completely untarnished. For instance, an ambitious project to install an innovative cell lysis process at the Mangere WWTP failed to deliver a serviceable asset, while poor initial scoping of the NSDWS upgrade project led to the final cost being several hundred percent over what had originally been budgeted in the AMP. But despite these setbacks the company still managed to deliver reductions in the prices of its water and wastewater services on the order of 20% in real terms between 2000 – 2008. This was at least partly due to strong internal control of capital and operational budgets. Indeed, a number of the company's Operation's staff took the time to grumble to the new CRM about the pressure they were under to keep Watercare's aging assets in service while also delivering year-on-year cost reductions. Such griping reveals, perhaps, a "view-from-within" that was not as rose-tinted as the view presented in the company's public documents. However, while there had been a few near-misses and close-calls, and although Watercare's Operations staff often grumbled about how the company was "operating on the edge"<sup>19</sup>, the company had never suffered a regionally significant interruption to service attributable to an asset, process, or management failure<sup>20</sup>. In other words, the company was "close to the edge" but had not yet "fallen over". This is a notable point, because, in terms of the Watercare's primary legislative

---

<sup>19</sup> One manager noted that some of Watercare's recent close-calls could easily have become major events resulting in widespread and very public service interruptions were it not for propitious timing and the swift responses and skills of Watercare's staff.

<sup>20</sup> The last major event was the 1993/94 drought which saw the enforcement of region-wide water restrictions. It was the region's subsequent reluctance to tolerate another such period of restriction which motivated Watercare to construct the Waikato pipeline and treatment plant.

objective, this is precisely where the company should have been. The “least-cost” service provision of the company’s legislative mandate required the company to balance the maintenance of service delivery (to required standards) against the long-term risk to the company’s assets (see the earlier sections of this appendix). While such a balance is very difficult to demonstrate objectively, Watercare’s successful service delivery track record, coupled with the long-term downward trend in the real costs of those services, can perhaps be contrasted with the discomfort expressed internally by Watercare’s Operations staff as a further indicator that the company was successfully achieving its core objectives. If the Operations staff were “comfortable”, with no gripes about the state of the company’s assets, then one might be tempted to suggest that the company was operating too far from “the edge” (i.e. gold-plating).

*Table I.7. Summary Watercare’s performance from 2000-2008*

Indicator	Trend 2000-2008	Inference
Non-compliance with regulatory standards (water and wastewater)	Consistent full compliance with NZDWS, except four minor breaches in 02/03 & 03/04.  Full compliance with resource consents at Mangere initially problematic due to limitations with the old plant and difficulties with the new plant after construction.  Significant improvement from 04/05 with only minor breaches caused by very high stormwater flows.	Operational performance reliably meets the required standards.
Variance against OPEX and CAPEX budgets	Opex variance is consistently within +5% to -4% of forecast budget.  Capex variance is consistently within 80-95% of forecast budget.	The company reliably achieves its Opex and Capex budgets (although it may be noted that the capital works forecasts are consistently optimistic).
WSAA benchmark score	Watercare recorded significant improvements in all but one category (which dipped slightly), with all 2008 benchmark scores between 70-85%.	The company’s asset management processes score well in the international benchmark test, and the increase in scores from 2004-2008 is indicative of a commitment to process improvement.
Drought security standard	With the exception of 2000-02, Watercare has consistently maintained compliance with the 1-in-200-year drought security standard. During 2000-02 the company was in the process of building and commissioning the Waikato treatment plant and pipeline to restore drought security.	The company consistently maintains a low level of drought risk through the timely investment in new infrastructure.

*Table I.7. Summary Watercare's performance from 2000-2008*

Indicator	Trend 2000-2008	Inference
MoH 'Aa' Grading	Consistent maintenance of the 'A' grading for the company's water treatment systems, and the 'a' grading for the bulk distribution network.	The company maintains a robust infrastructure supported by high quality of management systems and processes.
Credit Rating	Consistent maintenance of A or A+ credit rating.	The company's finances and debt servicing are well managed.
Lost Time Injury rate	The LTI reduced significantly over the period 2000-2008. Although it increased in 2008, it remains below the NZ Human Resources Benchmark Report benchmark of 7.	Watercare is successful at providing its employees with a safe work environment.
WW overflows not attributable to wet weather.	The percentage (by volume) of overflows from the wastewater collection system not attributable to wet weather flows is consistently in the range of 1 - 2% (varies from wet to dry years).	The wastewater system remains effective and reliable in terms of conveying dry weather flow volume.
Value-For-Money	Water prices remained virtually constant over the period 2000-2008, and, in fact, when adjusted for inflation actually reduced in real terms.  This is also the case for wastewater prices (the observed fluctuations are due to the volumes treated, i.e. from wet to dry years).	In light of the above statements of Watercare's performance, the significant reduction in costs, in real terms, on the order of 20% from 2000-2008 would seem to indicate that the company is consistently delivering Value-For-Money.

# Appendix II

## Early experiences crossing disciplinary boundaries

Appendix II explains the background to my study of Watercare's Corporate Risk Manager. There are three parts to that background:

- The original project on which I started by PhD research with Watercare;
- The practical work that I performed on various process-improvement projects at Watercare during the period of the initial research (2004 – 2006); and
- Framing a new inquiry following the termination of the original research project.

### **Early work: attempting to engineer sustainability**

The project on which I originally began my PhD candidacy (July 2004 to August 2006) was rather optimistically entitled “*Long-term planning for the sustainability of water service provision and infrastructure for the Auckland Region*” (Donnelly and Boyle 2004)<sup>21</sup>. It was motivated by a broad normative agenda to effect purposeful change in organisational behaviour and performance, where the ultimate objective was to achieve a “sustainable” Watercare organisation. The research embodied a radical agenda for change, with the underlying framework having strong similarities with other more well established

---

<sup>21</sup> The project was supervised under what was at the time, the Sustainability Engineering programme within the department of Civil & Environmental Engineering at the University of Auckland.

“sustainability” frameworks, such as The Natural Step (Robèrt et al. 2002), Critical Limits (Costanza and Daly 1987), or Critical Natural Capital (Ekins et al. 2003; Ekins and Simon 2003; Ekins 2003).

The logic behind that agenda was relatively simple. The ideal of “sustainability” holds that global environmental and social problems are a consequence of humans ignoring the fundamental dependence of human society on the ecological services provided by the natural environment (World Commission on Environment and Development 1987). Human actions (i.e. production and consumption) pollute, degrade, and destroy the natural environment, ultimately leading to the collapse of ecosystem functions and the destruction of that on which we depend for survival (Costanza 2006; Gunderson and Holling 2001; Scheffer et al. 2001; Steffen et al. 2004). It is believed by some that in a sustainable world production and consumption will, collectively, be in a form of dynamic balance with the natural environment (Costanza and Daly 1992; Ekins et al. 2003; Harris 2000). Since all significant forms of production and consumption are achieved through human organisations (formal, informal, and familial), achieving a “sustainable” human society, locally, regionally, and globally, therefore calls for substantial qualitative and quantitative changes in organisational behaviour and performance<sup>22</sup> (Lafferty and Meadowcroft 2000), although there is a broad spectrum of perspectives on the magnitude of the changes that will be required (Robinson 2004; Sneddon, Howarth, and Norgaard 2006).

The project plan drew inspiration from the Dutch Sustainable Technology Development programme (Weaver et al. 2000), but was fundamentally grounded in a functionalist systems paradigm and methodology typical of research and practice methodologies in engineering. The stated aim was to investigate alternative long-term strategic objectives for the future development of water supply and wastewater infrastructure in the Auckland Region. This was to be achieved by pursuing a research and modelling programme defined by the underlying “sustainability” framework (Boyle 2002, 2004b, 2004a)<sup>23</sup>. The original methodology involved five major steps: (1) future state visioning; (2) gap, risk and option assessment; (3) objective setting and gap analysis; (4) modelling; (5) risk identification and option finding. The project methodology was, ironically, an attempt to specify a rational comprehensive method of inquiry to determine, scientifically, what the conditions for the “sustainable” state

---

<sup>22</sup> I.e. reducing consumption of materials and energy, improving the efficiency of technologies, etc.

<sup>23</sup> Dr Boyle was the main supervisor of the original project within the Department of Civil & Environmental Engineering at the University of Auckland (additional sources for Boyle's framework and general thinking around sustainability are: Boyle 2004c, 2004d, 2005, 2007, 2008; Boyle and Coates 2005).



might be for water and wastewater infrastructure in the Auckland Region, which was intended to reveal the strategic changes needed to move development of the infrastructure in the “right” direction.

My work in 2004 and 2005 was primarily directed toward understanding the contestable global concept of sustainability and how to translate it into real-world practices. In this regard, I quickly became sceptical about the academic and practical validity of the proposed project methodology. There were two main components to my work. First, an exploration of “sustainability” literature, which became ever-expanding: sustainable development, sustainability and sustainability assessment; complexity, evolution, environmental science, and systems theory; problem framing and solving; scenario planning, land-use planning, and infrastructure planning and management; risk assessment; policy and evaluation theory; organisational management, governance, and decision-making theory; anthropology; socio-technological change, and the evaluation and management of technology; technological innovation, innovation systems theory, and social network theory; and, of course, the philosophy of intra- and intergenerational relations. This is illustrated in Figure II.1 (pg 379).

Second, concurrent with my exploration of the sustainability literature, I undertook to formulate a systematic conceptual framework for the translation of the world-level concept of “sustainability” into operational criteria in organisations. That problem ranged from the specification of the global environmental and social trends and moral judgements underpinning the imperative of sustainable development; to mapping the key scientific concepts and tools which provide understanding of the Earth system, and which support the calculation of environmental and social impacts and risks arising from organisational activities; to mapping the general operating context of organisations to understand the factors which enable and constrain the ability of organisations to “address sustainability”; to specification of the general requirements for “sustainability” for organisations. In this sense, the framework constituted an important analytical and sense-making device because it assembled and related the multitude of concepts that I encountered from my literature review toward a specific goal: defining guidelines for action, or rather for deciding about actions, in organisations. That framework therefore represented my early attempts to cross disciplinary boundaries.

I framed this work by conceptualising my role as that of a “sustainability expert” who might be contracted to advise organisations on how to approach the task of “achieving sustainability”. Thus, somewhat ironically in light of my later study of Watercare’s CRM, my guiding question was: what would a “sustainability expert” need to know? However, the fundamental dilemma that I encountered, and was eventually able to make explicit, was the

ultimately unbounded nature of the concept of “sustainability”.

Despite the degree to which the concept has permeated political and business discourse, it remains a vague term and concept marked by contention, contradiction, confusion, and conflict. While nearly everyone can agree that, at least in concept, it is a good idea (Jacobs 1999), no one seems to agree on what it means in practice. From one perspective, it means to incorporate environmental and social sensibilities into existing business paradigms, while from another nothing short of a radical shift in societal values, attitudes, and behaviours, including a whole new trajectory of industrial development is required for sustainability (Boyle 2004a; Hay 2002; Lafferty and Meadowcroft 2000; Peet 2000; Robinson 2004; Sneddon et al. 2006; Weaver et al. 2000). In essence, it seemed to me that just about anything could be claimed to be sustainable in one sense or another.

I eventually concluded “sustainability” is best understood as a socially constructed label which can be used in three very different, but equally legitimate ways (see Table II.1). But this realisation raised insurmountable problems with respect to my guiding research question (i.e. what would a “sustainability expert” need to know?). First, any given problem in society can be classified as a “sustainability” problem, consistent with the definition of sustainability adopted in any specific context. The science, knowledge and experience required to identify, frame, analyse, and solve the problem lie within some extant field of human endeavour, or, where a problem is multidisciplinary, across multiple fields. Thus, the bodies of theory, knowledge, and tools that a “sustainability expert” might need to call upon to identify, frame and solve a “sustainability” problem, are entirely dependent upon the specific issue in question and its context, and, collectively, could span the entire depth and breadth of human knowledge. And second, if this is the case, then the notion that “sustainability” can be a defined field of research or a professional body of knowledge in its own right is problematic at a fundamental level. Sustainability is simply a label which can be attached to any scientific endeavour consistent with the definition of sustainability adopted in any specific context, which would seem to suggest that a “sustainability science” (Kates et al. 2001) would ultimately be unbounded. Similarly, a generic “sustainability expert” is a problematic concept since this label can be attached to any professional depending on what problems they apply their knowledge and expertise to. A professional may offer advice under the guise of “sustainability”, but, fundamentally, he or she is acting as some other type of professional capacity (e.g. as an engineer, a business analyst, a risk manager, a planner etc.).

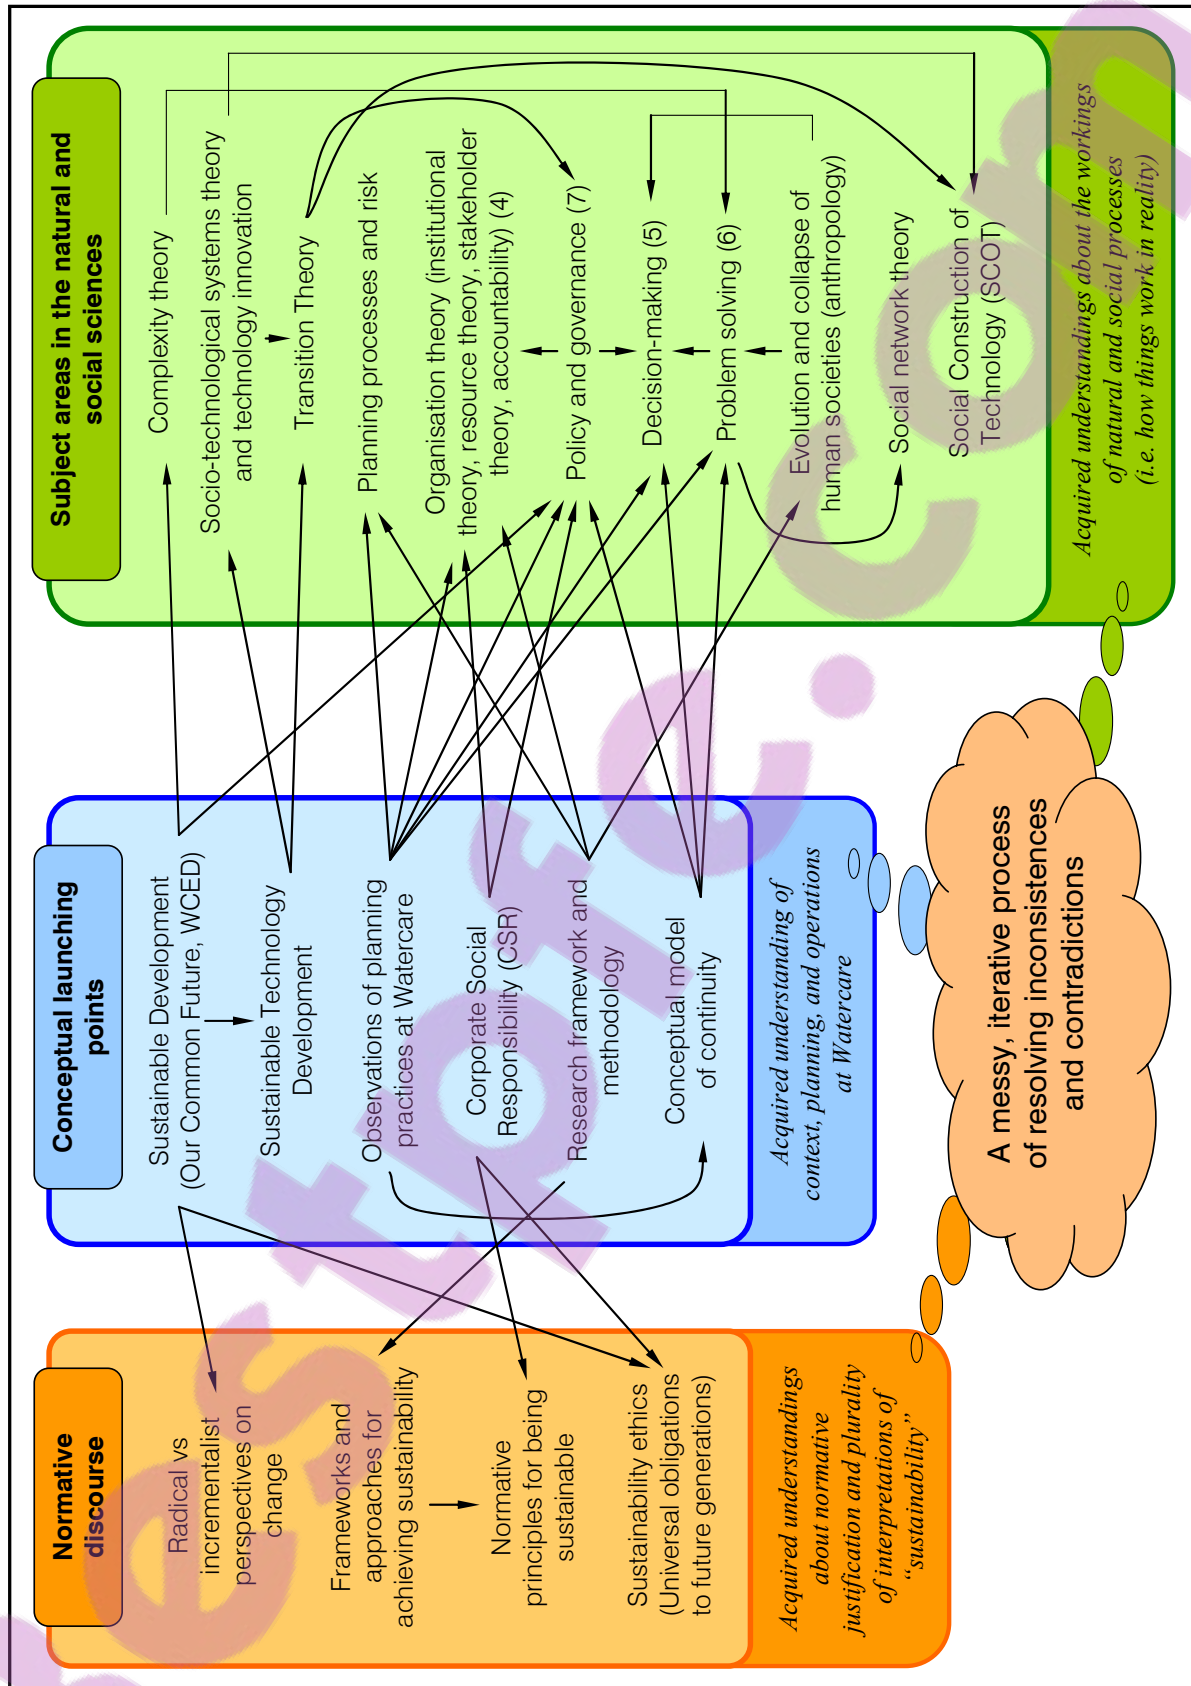


Figure II.1. Indicative conceptual map of my exploration of "sustainability" and related concepts. Arrows indicate linkages between concepts, numbers indicate density of connections to key subject areas in the literature.

**Table II.1.** Different usages of the “sustainability” label.

<p>Sustainability as:</p> <p><i>Continuity of physical and social processes</i></p>	<p>Here “sustainability” is a property of a process, function (process output), or outcome, being the ability of that process, function, or outcome to continue over a <i>definite</i> future period of time. Being “sustainable” means physical continuity.</p> <p>The fundamental practical issue with this conceptualisation lies with the problem of how to measure and predict the behaviour of complex physical and social systems in the future, and what constitutes a “sustainable” state in a dynamic environment.</p>
<p>Sustainability as:</p> <p><i>A future vision of social development</i></p>	<p>Here “sustainability” is a socially constructed vision of a future state of <i>human being</i> which is “sustainable”, where that future state is defined on the basis of normative and subjective ideas about how humans should relate to the natural environment and to each other. Such visions typically encompass elements of “Sustainability as Continuity”, above, and “Sustainability as Ethic”, below. Applied to an object, system, process, function, or outcome, “sustainability” in this sense refers to the degree to which the physical characteristics of these things comply with the system conditions derived from the normative vision. Being “sustainable” means to comply with what is deemed to be socially acceptable and desirable, for now and in the future.</p> <p>The fundamental practical issue with this conceptualisation lies with the problem of how to answer such inherently political questions as how the vision should be defined, who gets to decide on the relative merits of competing visions (ultimately all such visions are arbitrary), and how to enforce compliance with system conditions.</p>
<p>Sustainability as:</p> <p><i>An inter- and intra-generational ethic</i></p>	<p>Here “sustainability” is a universally applicable ethic defining the moral and human rights-based obligations upon the present generation with respect to current and future generations. Applied to an object, system, process, function, or outcome, “sustainability” in this sense refers to the degree to which those obligations are fulfilled. Being “sustainable” means to fulfill ethical obligations to future generations.</p> <p>The fundamental practical issue with this conceptualisation is how to measure “compliance” with moral obligations, particularly with respect to distant generations, and especially since compliance could be measured in a variety of ways the choice of which is made by the present generation (ultimately making the choice of measurement political).</p>

## Redesigning decision making infrastructure

During the period of the original research project (2004 – 2006) I had several opportunities to work on projects for internal process improvement, each of which was concerned with modifying some aspect of Watercare’s capital decision making infrastructure. At the time, I viewed these projects primarily as a means to gain exposure to and learn about Watercare and the company’s asset management and decision making processes. Much of the contextual information presented in Appendix I, and which informed the later reframing of my inquiry is a product of that exposure. However, reflecting on the objectives of those projects, and my approaches to them, revealed some themes which strongly parallel those

that emerged from my later shadowing of the Corporate Risk Manager.

## **Redesigning Watercare's risk assessment framework**

In early 2005 I reviewed the contents of Watercare's corporate risk register to develop an understanding of the company's risk profile. At that time the risk register was an MS Access database. The information held for each risk record included: a unique ID number, the risk classification (e.g. asset type and location), a description of the risk and causal factors, the inherent and residual risk assessment scores, and a description of the risk controls. After reviewing several hundred of the more than 1,500 risks in the database, I identified a number of issues that I felt could be improved upon, including (source: internal memo to Watercare's Risk Specialist, March 2005):

- Descriptions of risks and risk controls were often vague, non-specific, confused causes and consequences, or cited more than one causal factor;
- Risk assessments were qualitative. This meant that aggregation of individual risks to produce enterprise profiles was impossible;
- Risk assessments were static (i.e. were only indicative of present risk, not future risk);
- Sometimes the residual risk score did not reflect the specified controls;
- The use of the root-mean-squared averaging method was problematic;
- Records did not indicate what, if any, additional information and analyses were used to justify the risk assessment. There was no way to tell whether the assessment was the product of rigorous analysis or simply a casual guess by the assessor;
- Co-dependencies between risks were not accounted for;
- Enterprise risk profiles were reported simply as the number of risks in each class for each asset type.

The above issues are strikingly similar to those later identified by the new CRM and were clearly motivated by the same overarching concern with the objectivity of the information in the risk register. My expectation was that since risk was a representation of real potential future situations, the information in the risk register should be an objective (i.e. accurate and precise) assessment of the real risk faced by the organisation. I did not, at the time, consider how the information in the risk register might be different to what individual actors actually

understood about the real risk situation represented by that information.

In order to address some of the above issues I set out to redesign Watercare's risk assessment framework as a collaborative project with the company's Risk Specialist. We concentrated initially on a framework for evaluating risk in the company's main operational context (i.e. operating, maintaining, and developing the water and wastewater infrastructures; we were not concerned with the strategic or project management contexts for risk assessment). Our work addressed two questions: (1) which outcomes should be represented in the framework (i.e. which outcomes should be considered important)? And (2) how could those outcomes be represented?

### **Representing outcomes**

In the terminology of formalised risk assessment (see Appendix V and the discussion section of Chapter 4), risk represents uncertainty about future system performance, and should therefore be defined in terms of the performance objectives for a given system. From an Enterprise Risk Management perspective risk should be defined and measured in terms of the corporate objectives for the organisation. Thus, strictly speaking, outcomes are important when they result in the organisation failing to fulfil its performance targets or objectives, and the significance of those outcomes is measured by the degree of failure.

In 2005 Watercare's risk framework defined risk in terms of five consequence categories: Reputation, Finance, Environment and Public Health, Health and Safety, and Asset Management. Despite the generic and qualitative nature of the evaluation criteria used within the framework, there was clearly already some correspondence with the company's corporate objectives (i.e. the Sustainability policies and objectives defined in the Statement of Corporate Intent). The framework also reflected a common structure found in risk management standards and text books. We therefore did not consider the general scope of the existing framework to be in question. What mattered to us was the specific structure and level of detail in the framework.

However, reflecting on the nature of Watercare's operational context, it seemed that there would have been good reason to retain the breadth of scope in the existing risk framework even if Watercare's corporate objectives were not so broadly defined. As a bulk supplier, Watercare's formal market consisted of only six customers, the Local Network Operators, and that relationship was legally defined at certain physical points within the network by a mixture of regulatory and contractual specifications. But the infrastructure systems operated by each company were contiguous, forming an uninterrupted network

across the region. In this sense, Watercare's relationship with the LNOs, and indeed with the people and businesses of Auckland who were the final consumers of Watercare's services, was also clearly physical. What happened in Watercare's network, affected what happened in the LNO networks, and ultimately also the end consumer. Although changes in the performance of Watercare's network were generally not of sufficient magnitude to be noticeable to the end consumer under normal operating conditions (i.e. minor daily fluctuations in pressure and quality), this context required the company to maintain an extraordinarily high degree of reliability. Operational failures could have very real and very dire consequences far beyond the impacts to the company's reputation or financial statements. For example, a drought, earthquake or volcanic eruption, the collapse of a dam, equipment and process failure, damage to pipelines by third parties, and incidents in confined spaces, to name a few, were all easily imagined and understood as potential events which, depending on the situation, could result in serious consequences, not just for the company but also for Auckland's population, economy, and environment. Such consequences could include contamination or loss of the water supply, flooding of a sensitive ecosystem or popular beach with sewerage, illness, injury or death of employees or members of the public, damage to public infrastructure or private property, and in addition to the immediate and consequential financial costs, the prosecution of the company and employees, and loss of confidence in the company by its stakeholders. From a corporate social responsibility perspective (e.g. Carroll and Buchholtz 2000) there was clearly justification to define risk both in terms of the specific objectives of the organisation as well as more broadly in terms of outcomes generally considered to be undesirable on ethical grounds.

In regard to the selection of evaluation criteria, two factors were important. First, we wanted the framework to be simple to use. This was primarily a pragmatic criterion – if the framework were too complicated then no one would use it. It meant, however, that the criteria and the assessment method had to be easily understood by the majority of users. The second factor was objectivity. The performance implications of risk scenarios (e.g. equipment failure, asset deterioration, demand growth, changes in regulation) were usually well understood by engineers, planners, and managers. But the existing risk evaluation framework did not provide specific, contextually relevant criteria with which to represent those outcomes on paper. Rather, the qualitative criteria in the framework were so vague that, in practice, risk assessment became a largely subjective task. Not only was specificity lost in the translation from real outcomes into risk consequences, but this translation was subjective to the individual in question and therefore non-transparent and open to manipulation (even if these issues were addressed to some extent through discourse between the actors making

the risk assessment and those making investment decisions). These were the same concerns that would later motivate Watercare's CRM when he set out to redesign the corporate risk framework.

In re-designing the framework we wanted to make the evaluation criteria as concrete as possible, so as to reduce the amount of subjectivity involved in assessing risk. Drawing on the Risk Specialist's considerable industry experience we therefore looked for semi-quantitative indicators that reflected the real consequences of incidents, accidents, and failures in the operations context. This was a desktop exercise of imagining potential risk events (using the existing risk register for inspiration), and asking what the various impacts on those events might be (a) on the performance of the water and wastewater systems, (b) on the employees of Watercare, (c) on the environment and other 3<sup>rd</sup> parties (including Watercare's stakeholders), and (d) on Watercare, e.g. financial and legal impacts. In this regard, our approach paralleled that of the CRM later on. Unlike the CRM we did not try to explicitly draw risk criteria from the relevant performance standards, but instead relied upon the Risk Specialist's intricate knowledge of "what mattered", much as the CRM relied on experienced managers to identify the important criteria for his framework.

Our revised framework eventually employed eight main consequence categories: Water Service Delivery, Wastewater Service Delivery, Health & Safety, Damage to the Environment & 3<sup>rd</sup> Parties, Operational Impacts, Financial Impacts, Legal Liability, and Reputational impact.. In each category, outcomes were classified into various dependent and independent categories. So, for example, Water Service Delivery could be compromised in two ways, by a flow/pressure disruption and/or by a breach of quality standards, both of which were evaluated using three criteria, the significance of the disruption or beach, the duration of the incident, and the size of the population affected. For each evaluation criteria the upper (catastrophic) and lower (negligible) limits of severity were identified by asking: What is the worst that could happen? What is the least that could happen, with which the organisation should be concerned? The framework is summarised in Table II.2 (overpage)

To simplify the risk assessment and ensure consistency, we set up an MS Excel spreadsheet template which allowed users to simply select the relevant criteria from drop-down boxes. The software then calculated the consequence and risk scores. Since the consequence categories in our framework were not independent we initially opted not to aggregate the eight consequences scores into a single risk score (again this was something with which Watercare's CRM would later be concerned). The intention was that the eight scores represented the distribution (i.e. profile) of risk outcomes for any given risk event, which made explicit where the outcomes occurred and which ones were significant. Toward



the end of the development process, however, senior management asked us to include the calculation of a single risk score using the root-mean-squared averaging method that had been employed in the 2003 framework.

*Table II.2. Revised Risk Assessment Framework that I developed with Watercare's Risk Specialist in 2005*

<b>Categories (8)</b>	<b>Sub-categories (14)</b>	<b>Criteria (37)</b>
Service Delivery – Water	Flow / Pressure Disruption	Significance of disruption Duration of disruption Population (equivalent) affected
	Water Quality Breach	Significance of quality breach Duration of disruption Population (equivalent) affected
Service Delivery - Wastewater	Scale of Failure	Average flow in affected asset Type of overflows Duration of overflows
	Significance of Effects	Significance of effects Restrictions on public activity Duration of disruption
Health and Safety	Health and Safety	Fatalities Critical Injuries Serious Injuries Minor Injuries
Environment, 3 <sup>rd</sup> Party Damage	Environmental Impact	Significance of impacted environment Severity of impact
	Impact on Others	Scale of impact Severity of impact Duration
	Impact on Critical Service Infrastructure	Type of facility impacted Severity of damage
Watercare – Operational	Effect on Systems and Business Operations	System affected Scale of effect Magnitude of effects Duration of disruption to operations
	Impact on AMP / Programme of Works	Impact on forward works programme
	DWS Grading	Impact on DWS grading
Watercare – Financial	Financial Impact	Total repair and reinstatement costs Loss of Revenue Consequential losses borne by Watercare
Watercare – Legal Liability	Legal Liability	Breach of responsibility / regulatory breach Contractual breach with 3 <sup>rd</sup> party
Watercare - Reputation	Impact on Reputation	Visibility Community concern Loss of confidence in Watercare

## **Calibrating the new framework**

In order to test whether the new framework would produce reliable assessments in practice, a group of knowledgeable and experienced personnel were invited to use the framework to re-assess a range of existing operational risks selected from the risk register. The individual risks were assessed by the group, and then the positioning of each risk relative to the others was examined. The review considered whether the scores obtained were “representative of reality” with respect to both the particular event evaluated, and relative to the other events. Suggestions were made to adjust both the evaluation criteria and scoring system to get the resulting risk profiles and risk class boundaries “looking about right”.

Methodologically, the calibration exercise required the actors to judge whether the risk scores produced by the framework reasonably reflected their own perceptions of the significance of the events in question. It may be surprising that the new framework, so painstakingly crafted at a conceptual level to produce objective representations of risk, was subjected to a calibration against the supposedly subjective perceptions of individual actors. But this was the only way to validate the framework. It could not be compared against the results of the existing framework because the two were fundamentally different. It therefore had to be tested by exploring how well it represented Watercare’s “risk universe”. And, of course, the basis for comparison was not entirely subjective since the actors involved collectively possessed a considerable range and depth of experience in the water industry. In hindsight, however, I did not at the time appreciate the complexities or subtleties involved in formulating the second part of that inquiry, i.e. in eliciting and making explicit actors’ perceptions of reality in an objective manner. Even though the actors involved in the calibration worked for the same company, and were asked to consider a common reality, they necessarily approached that reality from quite different perspectives. A fundamental difficulty which arises is that these perspectives can and do lead actors to prioritise outcomes in very different ways. In this particular instance disputes over the relative scoring of different outcomes were relatively minor, but when the new CRM undertook essentially the same exercise with his revised framework in 2007 he encountered this problematic in a significant way (see Chapter 4).

## **Priorities and tolerance**

The process of re-designing the risk framework highlighted how such frameworks embed and make explicit the relative priorities of outcomes and objectives. The evaluation criteria and measures of severity in each consequence category were identified by considering the best

way to objectively represent and measure impacts. The scoring mechanism simply treated each consequence category as independent and generated a 0 – 100 score for each; the upper and lower bounds reflecting, respectively, the worst and least consequences that the organisation might reasonably be concerned with (at least in our opinion). However, this did not mean, for example, that the worst possible Health and Safety outcome could be directly compared with the worst possible Financial outcome. Such a comparison only becomes possible when the relative tolerance of the organisation to the outcomes in each category is known and made explicit. However, once such a determination is made and built into the framework, it essentially becomes locked in, i.e. Outcome A is always more important than Outcome B. In this regard, it was also apparent that while the framework might make explicit the relative priorities of outcomes and objectives, it could also efface the reasons why. The structure of Watercare's existing risk framework, for instance, implied that a score of say 75 in the Reputation category held the same relative level of significance for the organisation as a score of 75 in the Asset Management category. But the framework's supporting documentation was silent on how this relative priority was established, or whether the issue was even considered with the framework was originally developed.

Our development of the revised framework did not include an investigation of the relative significance of the different consequence categories. Formally, this was a task that could only be accomplished at Board level, and would therefore have to be addressed after the content of the framework had been confirmed. That is, one cannot reasonably ask the Board of Directors to evaluate their tolerance to risk outcomes if the categories for that evaluation are still uncertain. This points up the difficulty of changing formal decision making frameworks in bureaucracies. The new framework could only be adopted and used as a formal basis for decision making if the Board signed off on it. This would only happen if the majority of the senior management team and the Chief Executive agreed to recommend the revised framework to the Board. Consequently, changing formal decision making frameworks is not an easy task – considerable network building and enrolment of support is required.

By late 2006, thanks to the efforts of Watercare's Risk Specialist and Corporate Risk Manager (the predecessor to the subject of this study), the revised operations risk framework had acquired support at both senior management and Board level. However, the Corporate Risk Manager left Watercare at the end of 2006 and his replacement (and subject of the study in this thesis) took over in February 2007. Less than a month later the new Corporate Risk Manager announced that any further development of the risk framework was to be put on hold pending the development of an overarching vision for the future evolution of risk practices at Watercare.

## **Redesigning Watercare's capital planning framework**

The project to develop what was initially called the Planning Support Framework (PSF) was motivated by an independent assessment of the quality of Watercare's capital investment decision making process, which had been performed by an international management consultant in 2004 (prior to the start of my first research project). The report from that review recommended that Watercare make a number of improvements to its capital planning and decision making processes that should enable the company to better demonstrate a least cost approach (the least cost criterion is explained in Appendix I). These included improving the robustness of project estimates particularly in the first three years of the AMP, more robust application of cost-benefit analyses, incorporating social and environmental costs and benefits in those analyses, and developing the capability to demonstrate the total risk profile across the enterprise and how that profile will change over time with AMP expenditure.

Senior management at Watercare subsequently initiated a project to review and develop the company's capital decision making framework. The objective was to formalise a common and consistent basis for assessing and evaluating capital investment decisions across the enterprise. It was intended that the framework should: (i) provide an approach to expenditure evaluation that could be consistently applied across all levels and functions of the organisation, (ii) link evaluation criteria explicitly to business needs and objectives, (iii) in addition to risk and cost, allow the assessment of intangible (i.e. social and environmental) value, and (iv) provide a robust approach for assessing the impact of different levels of expenditure on enterprise performance over time, including the ability to test sensitivity to underlying assumptions and changing conditions (source: internal project documentation). The parallels with the objectives of Enterprise Risk Management are evident.

Since the idea of such an integrated evaluation framework was closely related to my research interest I was asked to undertake some development work on the project. I initially focussed on understanding the corporate objectives and the enterprise performance framework (i.e. that framework used to report enterprise performance in the Annual Report), the relationships between the company's planning and operations functions, the planning life cycle, and the existing basis for evaluating capital expenditure both for individual projects and for the AMP as a whole. I eventually concluded: (i) that the objectives and performance indicators specified in the SCI were inconsistent with those reported against in the annual report, (ii) that cost (net present value) and risk (defined by the framework described earlier) were the only consistent criteria on which investment decisions were evaluated, (iii) that

although the risk framework included social and environmental consequences, there was no consistent framework for evaluating the actual (rather than potential) impacts of investment decisions beyond financial cost, (iv) that formal methods of decision analysis (e.g. multi-criteria analysis) were only inconsistently applied and usually only on very large projects, and (v) that the qualitative nature of the risk framework prevented the aggregate analyses of risk profiles at the enterprise (i.e. AMP) level. Watercare's Corporate Risk Manager would later reach similar conclusions, which motivated his redeveloped of the corporate risk framework.

Consequently, at least in my opinion, the objectives of the PSF project could not be fully achieved within the company's existing capital investment evaluation capabilities. I assumed that new capabilities would have to be developed, starting with the specification of a comprehensive performance evaluation framework, the generic concept of which is shown in Figure II.2 (pg 390). My work on the project was subsequently focussed on defining at a conceptual level what an integrated performance framework might look like and the necessary relationships between that framework and Watercare's planning, operations, and strategic functions. My conceptualisation of those relationships is shown in Figure II.3 (pg 391). The main part of this work involved deriving a rationalised set of coherent and consistent evaluation criteria from the corporate objectives defined in the SCI and the Annual Report (both documents already included a number of objectives relating to the social and environmental performance of the firm). The objective was to define a set of performance indicators against which enterprise performance could be reported in the Annual Report, but which could also be used, with a minimum of manipulation, to evaluate individual capital investments at the project level (the centre box in Figure II.3).

I assumed that if such a set of enterprise-level indicators and project-level criteria could be specified, then this "common framework" would provide a consistent and robust basis on which to (i) evaluate the impact of both individual project expenditures and aggregated AMP investment profiles on the achievement of the corporate objectives going forward, and (ii) report the historical performance of the company in the Annual Report (left and right boxes in Figure II.3). I did not assume that it would ever be possible to accurately model and predict actual performance profiles taking into account every influence (internal or external) on the organisation. Rather, my assumption was that if a project were expected to alter the performance profile of the enterprise in any of the measured dimensions, and if this change could be estimated and specified quantitatively in terms of common performance indicators during the project evaluation, then it should be possible to generate an estimated enterprise profile ( $P_{t_2}$ ) at some future time ( $t_2$ ) for that performance dimension by combining the enterprise profile baseline ( $P_{t_1}$ ) and the expected change ( $\Delta P_p$ ) as a result of implementing the

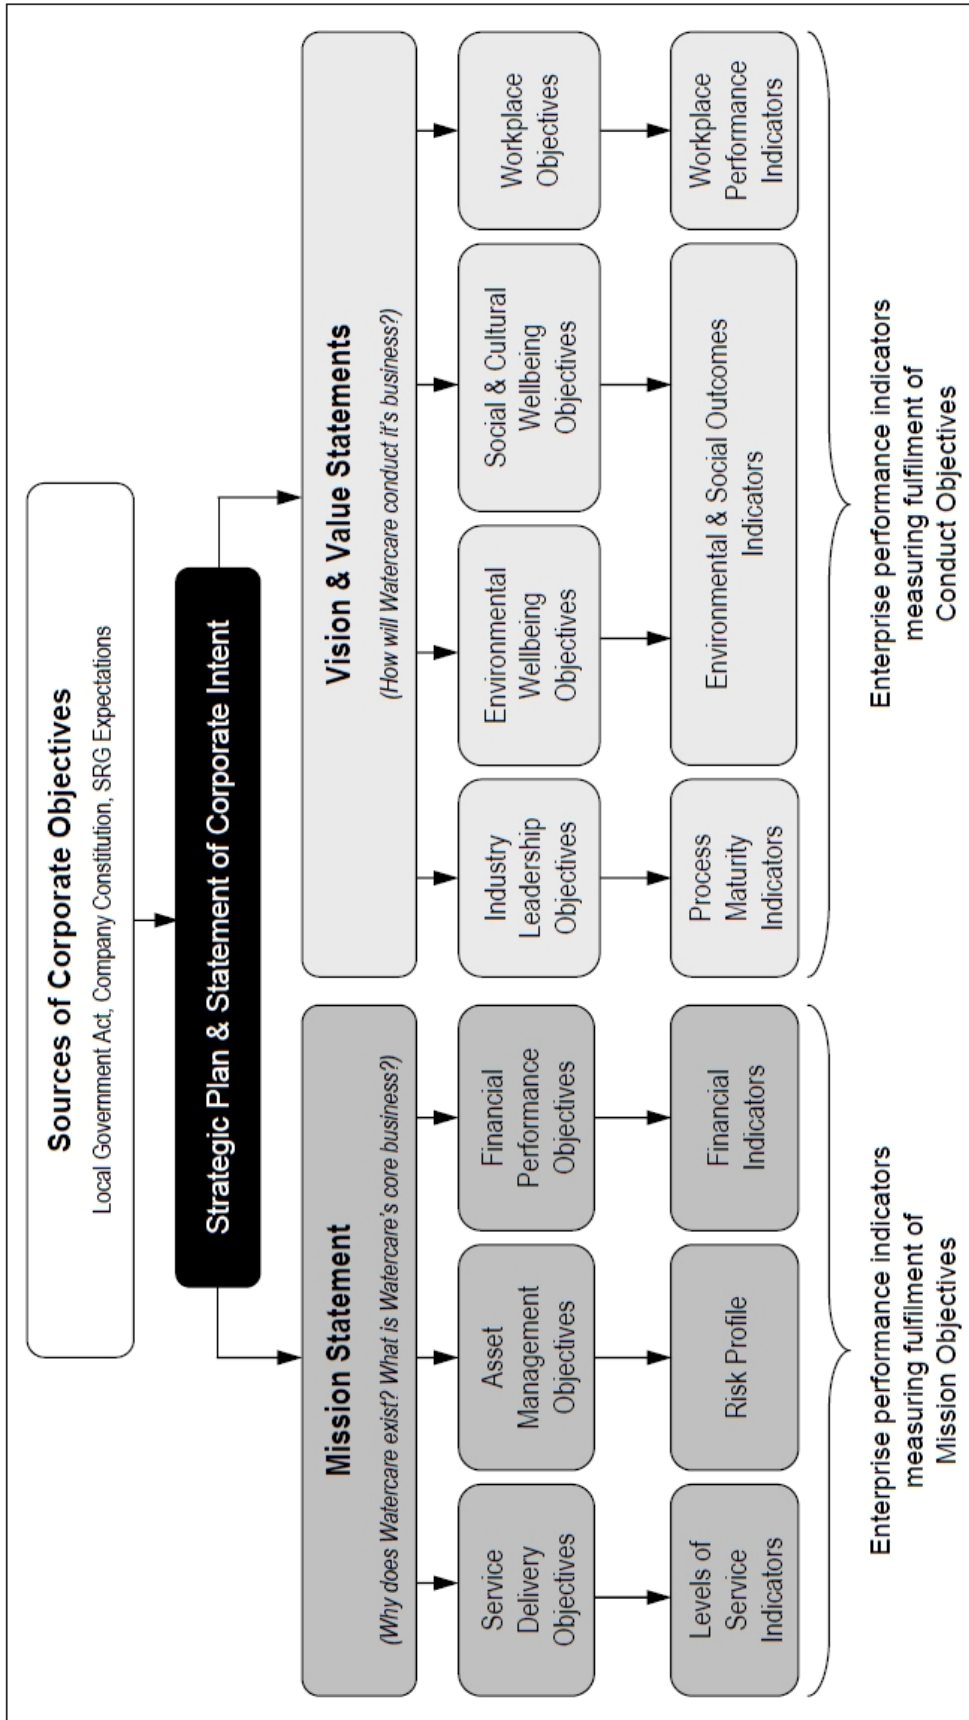


Figure II.2. Generic concept for translating Watercare’s corporate objectives into a comprehensive performance evaluation framework (adapted from presentation to managers, 26 May 2005).

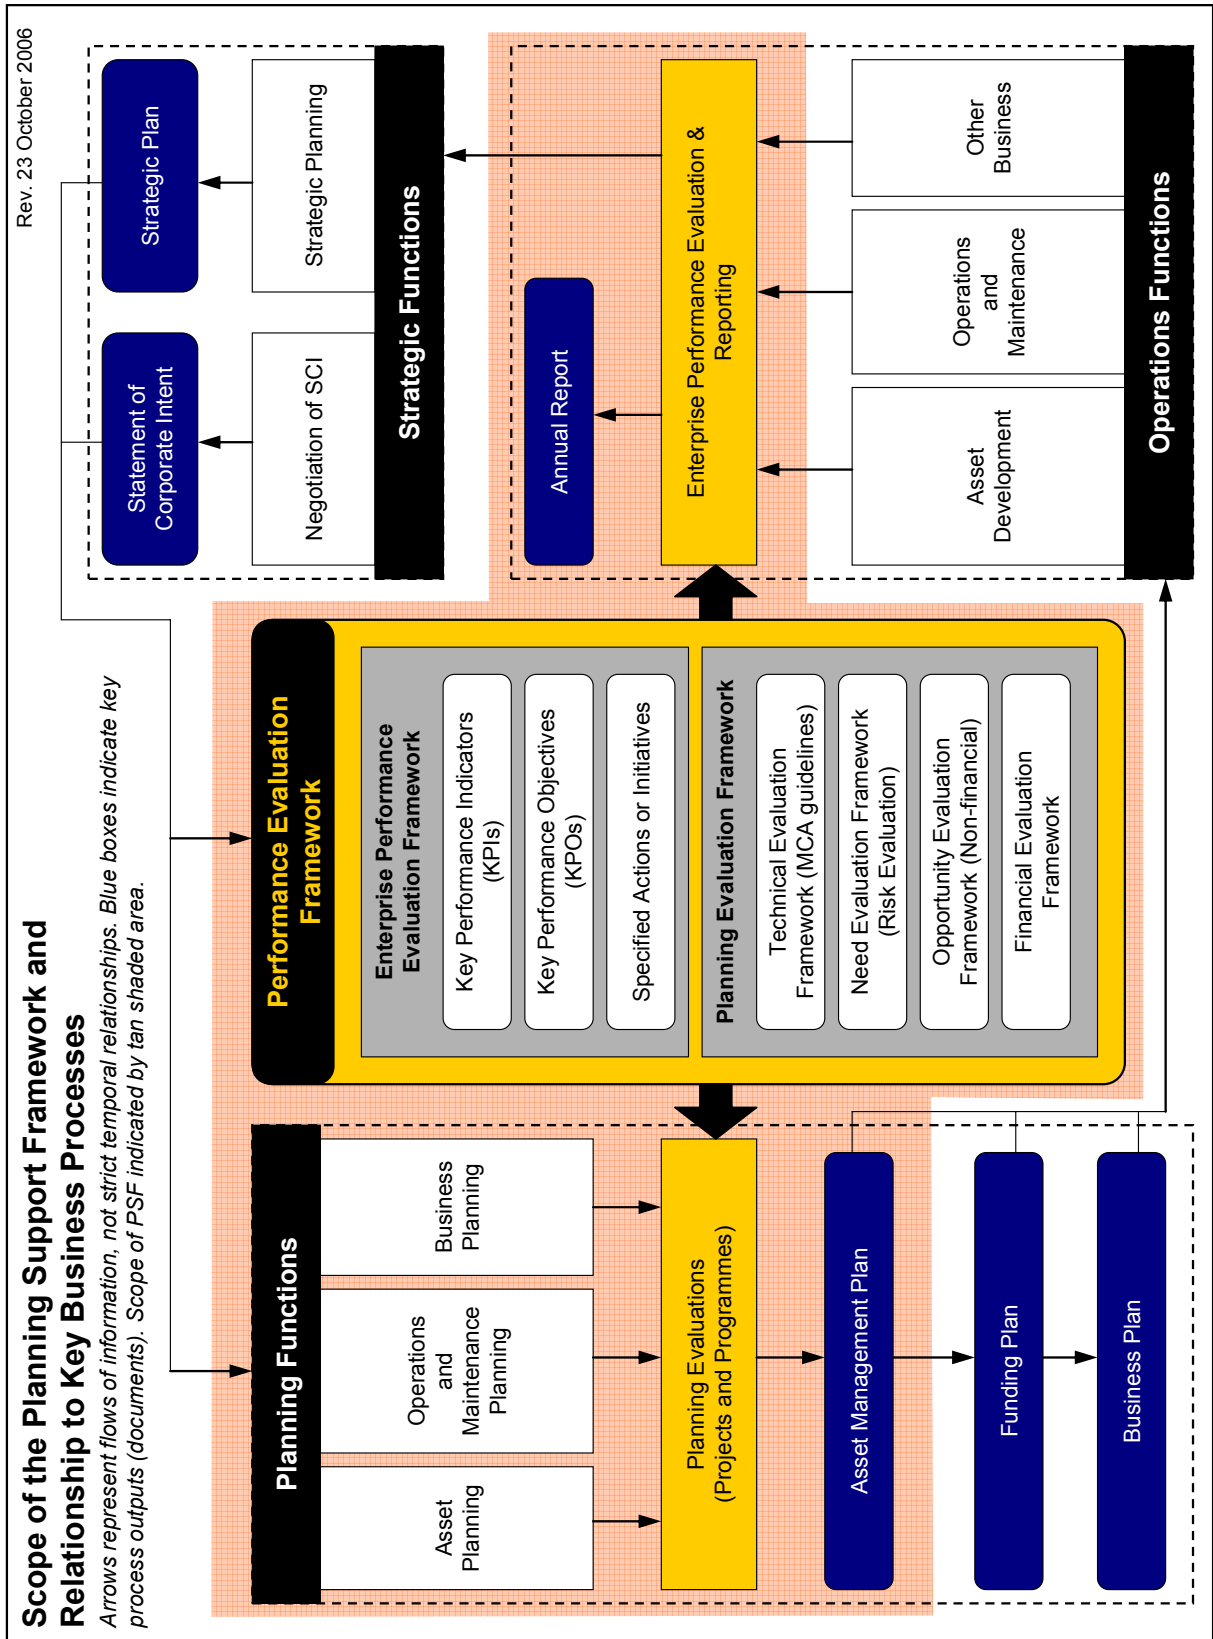


Figure II.3. Conceptual scope of integrated performance evaluation framework and relationships to key organisational processes (source: PEF, Proposed Solution, 29 November 2006)

project. That variable (P) would be represented as a step-wise function over time. Aggregating across all projects for each performance dimension (i.e.  $P_{t2} = P_{t1} + \sum \Delta P_p^{i-n}$ ), it should then be possible to estimate how the relative position of the organisation would vary across all performance profiles for different AMP scenarios, at least indicatively.

In this regard I was motivated and guided by several well established examples of similar frameworks used elsewhere, including: the project and programme evaluation framework used by the New Zealand transport funding agency (Transfund New Zealand 1997; LTNZ 2005), the UK Common Framework for Capital Maintenance Planning (UKWIR 2002), the International Benchmarking Network for Water and Sanitation Utilities Toolkit (IBNET 2004), the Water Services Association of Australia Asset Management Benchmarking Framework (WSAA 2003, 2008), the Baldrige Criteria for Performance Excellence (NIST 2006), and the Global Reporting Initiative Sustainability Reporting Guidelines (GRI 2002). The concept also allowed for the new evaluation framework to become a live process within the company, reflexively modified by the company's strategic functions to continually reflect the current objectives of the organisation (upper right hand box in Figure II.3).

By the latter half of 2006, the result of my efforts was a draft framework which specified a set of key performance indicators for the enterprise as a whole, and how they might be measured, and a corresponding set of evaluation criteria and procedures for capital projects specified across four independent categories (Economic evaluation, Risk evaluation, Technical evaluation, Intangibles evaluation). At that stage the conceptual work had been taken as far as it could go. Although there was still considerable development work to be done, I felt that I had at least demonstrated that the integrated evaluation framework was feasible (i.e. in principle there was no reason why the framework could not be put into practice). Further development would have to be practical, i.e. engaging with management, getting buy-in on the criteria and processes, and initiating a pilot study to develop the functional tools that would be necessary for measurement and calculation of the performance profiles. However, since I was not technically employed by Watercare, the project would have to be properly resourced if it were to proceed beyond the concept stage. Unfortunately, the period from late-2006 to mid-2007 saw the departure from Watercare of the three actors (the Corporate Risk Manager, the Chief Engineer, and a General Manager) who were the primary sponsors of the PSF project. This loss of support combined with the new agendas of the incoming personnel effectively put an end to the development of the project.



## Extending the Project Delivery Manual

Although the PSF project gradually ground to a halt by the end of 2006, it spawned a related project to extend Watercare's Project Delivery Manual (PDM). The PDM was an online, intranet-based manual which documented the processes, procedures, and tools for the detailed design and construction of asset projects (i.e. it applied to that phase of the asset life cycle labelled as "Project Delivery" in Figure I.3, Appendix I, p 352). The Operations group maintained a series of equivalent manuals documenting processes and procedures for the operation and maintenance of assets. There was, however, no comparable documentation for the work undertaken by the Planning group (i.e. for those stages of the asset life cycle labelled as "Project Definition" and "Project Development" in Figure I.3).

The Planning Support Framework had included at least notional definition of Watercare's business processes, principally as a means to identify key points where the framework would be used for various evaluation purposes (e.g. Figure II.3, above). As a spin off from the PSF project I was asked to extend that work to produce the necessary documentation as an addition to the front-end of the existing PDM. This work involved meeting with key members of the planning group to learn the development life cycle of a project from the identification of the initial business trigger through to the hand over of the scheme design to the project delivery team, and then documenting that process in the form of a series of generic process diagrams and accompanying explanatory text.

The work was not, however, simply descriptive, it was also a task of formalisation. The investigative process revealed that there was a broadly understood "project lifecycle" that, for each project, produced three primary administrative documents – the project information sheet and the 20-year operational and capital expenditure estimates for the project (published in the Asset Management Plan), and the formal request for capital expenditure (the CapEx, produced for management signoff before the project moved into the delivery phase). Beyond these three documents, additional project information and documentation was filed in a variety of locations, including the official project file (which utilised a universal coding scheme), in a folder (or folders) on the planning computer network, in Watercare's formal electronic archiving system, and in various filing cabinets within the office. Although formal procedures for the filing and storage of project documents did exist, they were not consistently used by members of the planning team. This produced some inefficiencies with respect to knowledge retention and transfer regarding existing projects. There was anecdotal evidence, for example, that it was possible for two similar or overlapping projects to proceed through the planning process independently because their respective planners were unaware

of the existence of the other project, or that investigative work was sometimes repeated because the planner was unaware that it had already been done on an earlier project.

The problem was not so much the loss of project information and data, but rather the forgetting, over time, that it existed at all and of where it was physically located. In an effort to resolve some of these inefficiencies I included in the documentation of the project development lifecycle certain standard forms and reports that were intended to capture key information about the project, what additional information and documentation existed, and where that information was stored. My intention was that completion of these forms and reports would be a compulsory requirement for management sign-off of the project, and would be held in a 'central projects register' (which would have to be developed) accessible via the company's intranet, and which would form a central repository of up-to-date information for all past and present projects cross-referenced against the company's infrastructure assets. The idea was that the database would become a single point of entry to the universe of projects investigated and constructed by the company.

## **More transdisciplinary work: reframing the inquiry**

The original research project, described above, was effectively terminated by August 2006 as a result of a breakdown of the student-supervisor relationship within the university caused by entrenched disagreement over how to approach the task of operationalising a concept which clearly transcended the boundaries of the engineering profession. At that time I was unwilling to abandon my PhD without first investigating whether I could complete a project by building on what I had already done. I had retained my privileged access to the Watercare organisation, and also the project funding. That funding was conditional on the research seeking a practical output in the form of a technological development, or improvements to business planning processes around the firm's core technology programme. I was therefore in the relatively unique position of having research funding and an organisational setting in which to conduct research, but no project. The following sections present the key conceptual work that I undertook to identify what that project should be.

I had also retained, out of the original project, my general interest in what I referred to as "change processes in Large Technical Systems". The reframing process was therefore one of drawing together concepts from relevant literatures to frame up my understandings of Watercare's operations and context to constitute, in essence, a general systems model of technological change in Watercare's infrastructure. The following sections do not present a

precise and definitive account of the Watercare system, but rather develop a range of concepts which are intended only to serve a sensitising function. The distinction is as follows:

A definitive concept refers precisely to what is common to a class of objects, by the aid of a clear definition in terms of attributes or fixed benchmarks. This definition, or the bench marks, serve as a means of clearly identifying the individual instance of the class and the make-up of that instance that is covered by the concept. A sensitizing concept lacks such specification of attributes or bench marks and consequently it does not enable the user to move directly to the instance and its relevant content. Instead, it gives the user a general sense of reference and guidance in approaching the empirical instances. Whereas definitive concepts provide prescriptions of what to see, sensitizing concepts merely suggest directions along which to look.... [Sensitizing concepts] rest on a general sense of what is relevant. (*extract from Blumer 1969, pp 147-148, quoted in Kovach 2004, p 38*)

The task of reframing the research was one of structuring a context-specific understanding of “what is relevant” with respect to the evolution of the infrastructure as a basis on which to suggest “directions along which to look” for opportunities for viable inquiry. In this regard, the following sections do not frame up a specific inquiry, but merely answer the following question: “If I want to understand how change occurs in this system, then where should I look, and to what should I pay attention?”

The reader is referred to Appendix I for a description of the Watercare organisation, its infrastructure, institutional context, and asset management processes. The first section, below, considers Watercare as an agent within a Large Technical System and characterises the processes of technological change within that system. The second section draws attention to Watercare’s business planning and engineering design functions as the key processes of technological change for the infrastructure at large. The next section elaborates the analytical and creative nature of the engineering design process, since design was a core component of Watercare’s planning capabilities. The final section draws on the earlier material to characterise Watercare’s planning and design processes as the company’s primary dynamic capabilities, and, subsequently, to frame up contexts of potential inquiry.

## **The evolution of Large Technical Systems**

Within science and technology studies (STS) there is a domain of inquiry which investigates the development and evolution of Large Technical Systems (LTS) and the societal implications of infrastructural changes (van der Vleuten 2004). Although there is no consensus on a strict definition of LTS (van der Vleuten 2004), researchers in the field have studied a variety of society-wide infrastructures, including electricity supply networks

(Hughes 1983; Levy-Leboyer 1988), railroad networks (Salsbury 1988; Caron 1988; Heinze and Kill 1988), telecommunications and the internet (Galambos 1988; Bertho-Lavenir 1988; Thomas 1988; Davies 1996; Abbate 1994), air traffic control (La Porte 1988), nuclear power (Rochlin 1994), water supply systems (Katko, Juuti, and Pietila 2006; Monstadt and Naumann 2005; van der Brugge, Rotmans, and Loorbach 2005) and highways (Geels 2007); among others.

These large infrastructure systems are highly complex, geographically dispersed networks of heterogeneous technical and non-technical components (Hughes 1987; Joerges 1988; Monstadt and Naumann 2005; van der Vleuten 2004). While the network form of the technical artefacts is obvious, the definition of LTS also encompasses the people and organisations which develop, operate, and maintain the physical infrastructure, the scientific (knowledge) artefacts that they employ, and the institutional arrangements (cultural, political, legal, and economic) in which they are embedded. These functional social systems have their own unique knowledge bases, particular norms, specialised occupations, large formal organisations, and intra- and inter-organisational (social) networks (Monstadt and Naumann 2005). In this sense, technology is defined broadly, encompassing both physical 'hardware' and social 'software', "[s]oftware represents the knowledge, know-how, practices and organisational skills needed to develop technologies and to produce and use artefacts...[and] is reflected in organisational and institutional arrangements..." (Weaver et al. 2000, p 46).

Following his studies of the emergence and development of electrical power networks in Europe and the U.S. from 1880 to 1930, Hughes (1983, 1987) was the first to offer a description of the evolutionary process by which LTS emerge and develop from small local-scale systems into large integrated regional and national-scale systems. He suggested that LTS develop through a number of distinct phases identifiable by the primary type of activity occurring during that phase: initially invention, development, and innovation, and then transfer, growth, competition, and consolidation (Hughes 1987; Joerges 1988). Similar patterns of evolution have been described elsewhere: Weaver et al. (2000), for example, described technological change in society as a three-phase process of invention, innovation, and diffusion; Gökalp (1992) suggested a four-phase process, the initial phase, the accelerated development phase, the stabilisation phase, and the decline phase; while Rotmans et al. (2001), van der Brugge et al. (2005), Geels (2007), and others of the Dutch school have described transitions between dominant technological forms in large socio-technical systems also in four phases – pre-development, take-off, acceleration, and stabilisation. Hughes (1987) argued that the evolutionary process should be understood as a pattern rather than a model

– the process is not linear, the phases are not simply sequential or independent, they may not necessarily occur in the order presented below, and iterative loops between phases may be identifiable in the history of a particular system.

### **The initial phase: invention, development, innovation**

In the initial phase, a new technological system is invented which is radically different to existing forms of technology, and which constitutes a new solution to an existing societal problem (Hughes 1987; Joerges 1988; Weaver et al. 2000). Contrary to popular belief, new technologies are rarely the products of accidents or moments of inspired creativity by lone inventors; rather, creative inventiveness is situated in particular social and historical contexts from which it cannot be separated (Latour 1987, 1988; Weaver et al. 2000). The very possibility of invention (whether intentional or accidental) requires the bringing together and arrangement, in a certain time and place, of particular knowledges, know-how, materials, and technological artefacts, all of which takes considerable time, effort, and resources to accomplish (Latour 1987, 1988). Further, this ‘bringing together’ is accomplished within broader contexts of social relationships, cultural and political values, organisational and institutional arrangements, and supporting infrastructures, and of certain problems and opportunities, intentions, expectations, incentives, and competing interests, all of which support or resist the inventive process (Latour 1987, 1988; Pickering 1993, 1994).

Once invented, the success of a technology is never guaranteed, no matter how revolutionary or novel – the uptake of technology in society is not simply a function of some inherent value of that technology. In order to survive, the new technology must be made practical and efficient (innovation), and must be provided with the financial, political, and legal support necessary to move it from the laboratory to the market place (development) (Hughes 1987; Joerges 1988; Latour 1987, 1988; Weaver et al. 2000). Niche markets play an important role here (de Bruijn and Norberg-Bohm 2005). This is a difficult and costly process of network building, accomplished in Hughes’ original studies by entrepreneurial inventor-engineers, or “system builders”, who both created and linked together a vast number of heterogeneous elements, technical and non-technical, to create and develop their systems (Hughes 1983, 1987; Joerges 1988). They were not just technologists, but were entrepreneurial in multiple senses, technological, economic, and political (Joerges 1988). For example, to develop his system for domestic electrical lighting, Thomas Edison...

... had to establish not only ties among investors, politicians, and technicians, but also circuits for the transmission of capital into his enterprises, generating stations to transform coal into

electric power, carbon filaments whose resistance was calibrated to the current-carrying capacity of copper cables and to the cost of the copper, a system of patents and the means to enforce them, and cable networks to carry direct or alternating current from place to place (Hughes 1983)... The Edison Electric Light Company, set up in 1878, did not sell lighting. It held patents on the devices Edison's team invented—light bulbs, generators, distribution systems—and licensed or sold the patents around the world to raise income and attract investment to finance Edison's workshops, experiments, and demonstration projects (Hughes, 1983, p. 39). It organized capital flows through networks of lawyers, legislation, patent enforcement, and publicity. Edison's first central generating station began commercial operation in 1882 in New York City close to Wall Street. The location was chosen to attract the attention of financiers, and because the half-mile radius its distribution network could reach included many shops and restaurants, which would draw customers and publicize the system (Hughes, 1983, p. 41).

*(excerpt from Mitchell 2008, p 1117)*

### **The latter phase: transfer, growth, competition, consolidation**

Following the initial development phase, if the technology is successful then a pattern of broader adoption and diffusion within society becomes evident (Hughes 1987; Weaver et al. 2000). The technology is transferred to new locations, where it must be adapted to environments (i.e. physical, political, legal, and cultural conditions) different to that under which it was originally developed (Joerges 1988). This is followed by the accelerated growth and development of the system, not just in terms of the technical hardware, but also reflected in an accumulation of socio-cultural, economic, organisational, and political and legal institutional changes – the growth and development of the supporting technological software (Rotmans et al. 2001; van der Brugge et al. 2005; Weaver et al. 2000). These latter stages are also marked by competition, where rival technological systems may emerge and compete for dominance, and consolidation, where rationalisation, efficiency, and capital intensification become the dominant system goals, and geographically and organisationally fragmented systems gradually become integrated into large regional and national-scale networks (Hughes 1987; Joerges 1988).

Hughes used the term “momentum” to capture the notions of a build-up of a giant mass of innumerable technical and organisational components, and of velocity, in the sense of expansiveness, rate of growth, and goal-directedness (Hughes 1987; Joerges 1988). Even from early beginnings, LTS exhibit path-dependency, where strategic investment decisions orient the technological development of the system in certain directions, and once set on those paths, the sheer mass of technological and organisational system components creates a kind of developmental momentum which inhibits flexibility and adaptability (Hughes 1987; Joerges 1988; Katko et al. 2006; Monstadt and Naumann 2005). LTS, particularly in their

mature stages, are also characterised by high capital intensity, long planning and pay back periods, high sunk costs of investment, and (usually) long and complex authorisation procedures, which further contribute to inertia against change (Monstadt and Naumann 2005).

However, throughout the evolution of the system, the technology (both hardware and software) is subject to what Hughes' called "reverse salients", understood as technical or organisational anomalies resulting from the uneven elaboration or evolution of the system (Joerges 1988). In essence, they are technical, organisational, or social problems encountered as the system grows and becomes more complex, and are essential to the evolutionary process because, in requiring solutions to critical underlying problems, they drive the continued inventive activity and growth of the system (Joerges 1988). Such problems give rise to two possible types of innovation (Joerges 1988; Weaver et al. 2000). Conservative or incremental innovation refers to solutions which optimise the technical and economic performance of the existing system. This type of gradual improvement is driven by competitive pressures and is facilitated by the learning process that occurs through cumulative experience in producing and marketing a particular technology (Weaver et al. 2000). In contrast, Schumpeterian innovation refers to radical or trend-breaking solutions which take the form of an entirely new process, product, or organisation, and which may eventually give rise to new competing technological systems:

The new breakthrough will have a genealogy traceable to earlier discoveries. Nonetheless, the essence of a radical breakthrough is that it represents a fundamentally new approach that departs from pre-existing engineering practice and technologies and is not a continuous development of any single former approach. (*excerpt from Weaver et al. 2000, p 49*)

A paradigmatic example of radical innovation is the fluorescent lamp, which represented a breakthrough solution to the problem of the energy efficiency and cost of incandescent bulbs for high-intensity lighting (Weaver et al. 2000).

### **Transitions in LTS**

Despite their developmental momentum and inertia to change, LTS can and do undergo transitions from one dominant arrangement of technical and social components to another (Geels 2007; Rotmans et al. 2001; van der Brugge et al. 2005). This notion of transition is understood as a gradual process of change not just in the technical components, but also the emergence of new linkages, new knowledges, new rules and regulations, new organisations and institutions, changing roles for existing actors, and new patterns of production and

consumption (Geels 2007; Kemp, Parto, and Gibson 2005; Rotmans et al. 2001). Historical examples include the transition from 'horse and carriage' to the private automobile, from mainframe computing in the 1970s to desktop computers in the 1990s, and from prop-driven to jet-propelled aircraft in the airline industry. In this sense, the evolution of LTS from new technology to mature infrastructure system represents a series of complex, non-linear transformations in the socio-technical landscape involving a broad heterogeneity of elements and actors across multiple scales, and historically occurring over periods from a few to several decades (Geels 2007; Rotmans et al. 2001; van der Brugge et al. 2005).

At the niche or micro-level, individual actors (people, companies) and local practices are distinguished – it is at this level that new ideas and initiatives, innovations, new techniques, alternative technologies, and social practices emerge. The regime or meso-level refers to broader networks of established organisations and institutions, and it is here that the dominant technologies, practices, rules, norms, and policies of specific socio-technological regimes are found which serve to both stabilise and optimise the system. The macro-level refers to the overall societal landscape in which processes of change occur, consisting of cultural and political values and worldviews, political coalitions, the natural and built environments, and political, economic, and legal institutions.

Existing LTS are defined and sustained by arrangements of technical and social components at the meso- and macro-levels, while new possibilities for development emerge at the niche level in the form of innovative breakthroughs in technologies and practices (Geels 2007; Rotmans et al. 2001; van der Brugge et al. 2005). Depending on the network-building skills of associated actors, these innovations may more or less successfully evolve into competing socio-technical systems, thus sustaining the process of transition (Hughes 1987; Rotmans et al. 2001; Weaver et al. 2000).

### **Distributed control in complex LTS**

Hughes conceptualised LTS as complex problem-solving systems, the problems having “to do mostly with reordering the physical world in ways considered useful or desirable, at least by those designing or employing” the system (Hughes 1987, p 53). In this sense, LTS are seen as systemic solutions to significant societal problems.

From the foregoing, it should be clear that large technical systems may take a variety of forms, in terms of the complexity and dynamism of the technical and social components. As a simple heuristic, complexity refers to the number and heterogeneity of components in the system, while dynamism refers to the rate of change in those components over time: systems



with a large number of dissimilar components are more complex than systems with a lesser number of similar components; and systems where the components are in a continual state of flux are more dynamic than systems where the components change only infrequently (Duncan 1979).

A strict application of these concepts is not necessary. The point is that large technical systems may be more or less complex and dynamic, but the more complex and dynamic the system, the greater the degree to which problem solving in that system must be seen as a distributed, emergent phenomenon (Diamond 2005; Tainter 1988). At the meso- and macro-levels, changes in the system cannot be understood as supposedly rational solutions to problems identified and analysed by “the system”. The system does not make decisions, as that term is commonly understood. Rather, at these levels, observable changes are emergent effects of heterogeneous, interacting problem solving activities carried out by individuals and organisations within the system, but not directly or causally attributable to any single one. The transition models developed by the Dutch school are explicit about this – macro-level socio-technical transformations are emergent effects of heterogeneous interactions within a multi-level structure of nested systems (Voß and Kemp 2006; there is a strong parallel here to theories of dynamic change in ecological systems, see for example Gunderson and Holling 2001; Lansing 2003; Waldrop 1992).

It follows that the more complex and dynamic the system, the less any single individual, organisation, or institution can exert control over the system. Governance in such systems is distributed; where governance is understood as the capacity to influence the activities and directions of development in the system at large (Voß and Kemp 2006). Capacities to influence the performance and technological orientation of large technical systems are located, to varying degrees, with a broad range of actors, including central and local government, regulators, interest groups, commercial and public sector organisations, producers and consumers, scientists and engineers, and the media (Voß and Kemp 2006). It further follows that efforts to facilitate directed change in LTS must be grounded in understandings of evolutionary change processes at the niche, meso, and macro levels (e.g. the notion of transition management, Kemp et al. 2005). De Bruijn and Norberg-Bohm (2005), for example, report a number of empirical case studies from the United States and Europe of efforts to facilitate industrial transformation through approaches both at the industry/sector level and at the firm level. Weaver et al. (2000) report on the Dutch experience with targeted programmes to build systemic capacities for technological innovation in large technical systems in The Netherlands.

## Characteristics of change in Auckland's water infrastructure

Watercare was a primary management organisation within what may be generally termed a Large Technical System (Bijker, Hughes, and Pinch 1987; Mayntz and Hughes 1988). Watercare's core business was the development, operation, and maintenance of a complex, geographically extensive infrastructure system supplying bulk water and wastewater services to around one million people in seven cities and districts in the Greater Auckland Region (see Appendix I for further detail on the company). With the exception of a few minor water supply and wastewater schemes serving some small rural townships, Watercare owned and operated all of the raw water sources, water treatment plants, and associated bulk supply pipelines in the region, and the larger of the two major wastewater treatment plants. This infrastructure system, together with the organisations that constituted the Auckland Water Industry and various institutional mechanisms (e.g. Acts of government, regulatory standards, local governance structures), constituted a socio-technical solution to the water-related aspects of the societal problem of protecting public health in the Auckland Region.

The literature on LTS emphasises the distributed nature of individual actors' abilities to influence control over the system at large. In the absence of a comprehensive hierarchy of control, transitions from one dominant technological form to another are characterised as involving the emergence and diffusion of new technical and organisational arrangements across a range of scales (Geels 2007; Hughes 1987; Rotmans et al. 2001; van der Brugge et al. 2005; Weaver et al. 2000).

Examination of Watercare and the Auckland Water Industry revealed a somewhat different picture of technological evolution. Figure II.4 presents an indicative mapping of actors and the forces they actually or potentially exerted on the region's water and wastewater infrastructure. An important feature of the local system was that control of the infrastructure was compartmentalised rather than distributed. Direct control of the planning, development, operation, and maintenance of Auckland's water and wastewater infrastructure lay only with Watercare and the six Local Network Operators, each of which was responsible for independent but contiguous sections of the network (represented by the dotted boundary in Figure II.4). The other actors in Figure II.4 exerted various drivers and constraints on the system, but their influence was only indirect. The physical infrastructure did not change unless and until Watercare and the LNOs took action. This is not to say that Watercare and LNOs could ignore the forces exerted by the other actors in the system. Rather, what mattered was that Watercare and LNOs were the only actors able to physically change the infrastructure, and their spheres of control, while contiguous, did not overlap.

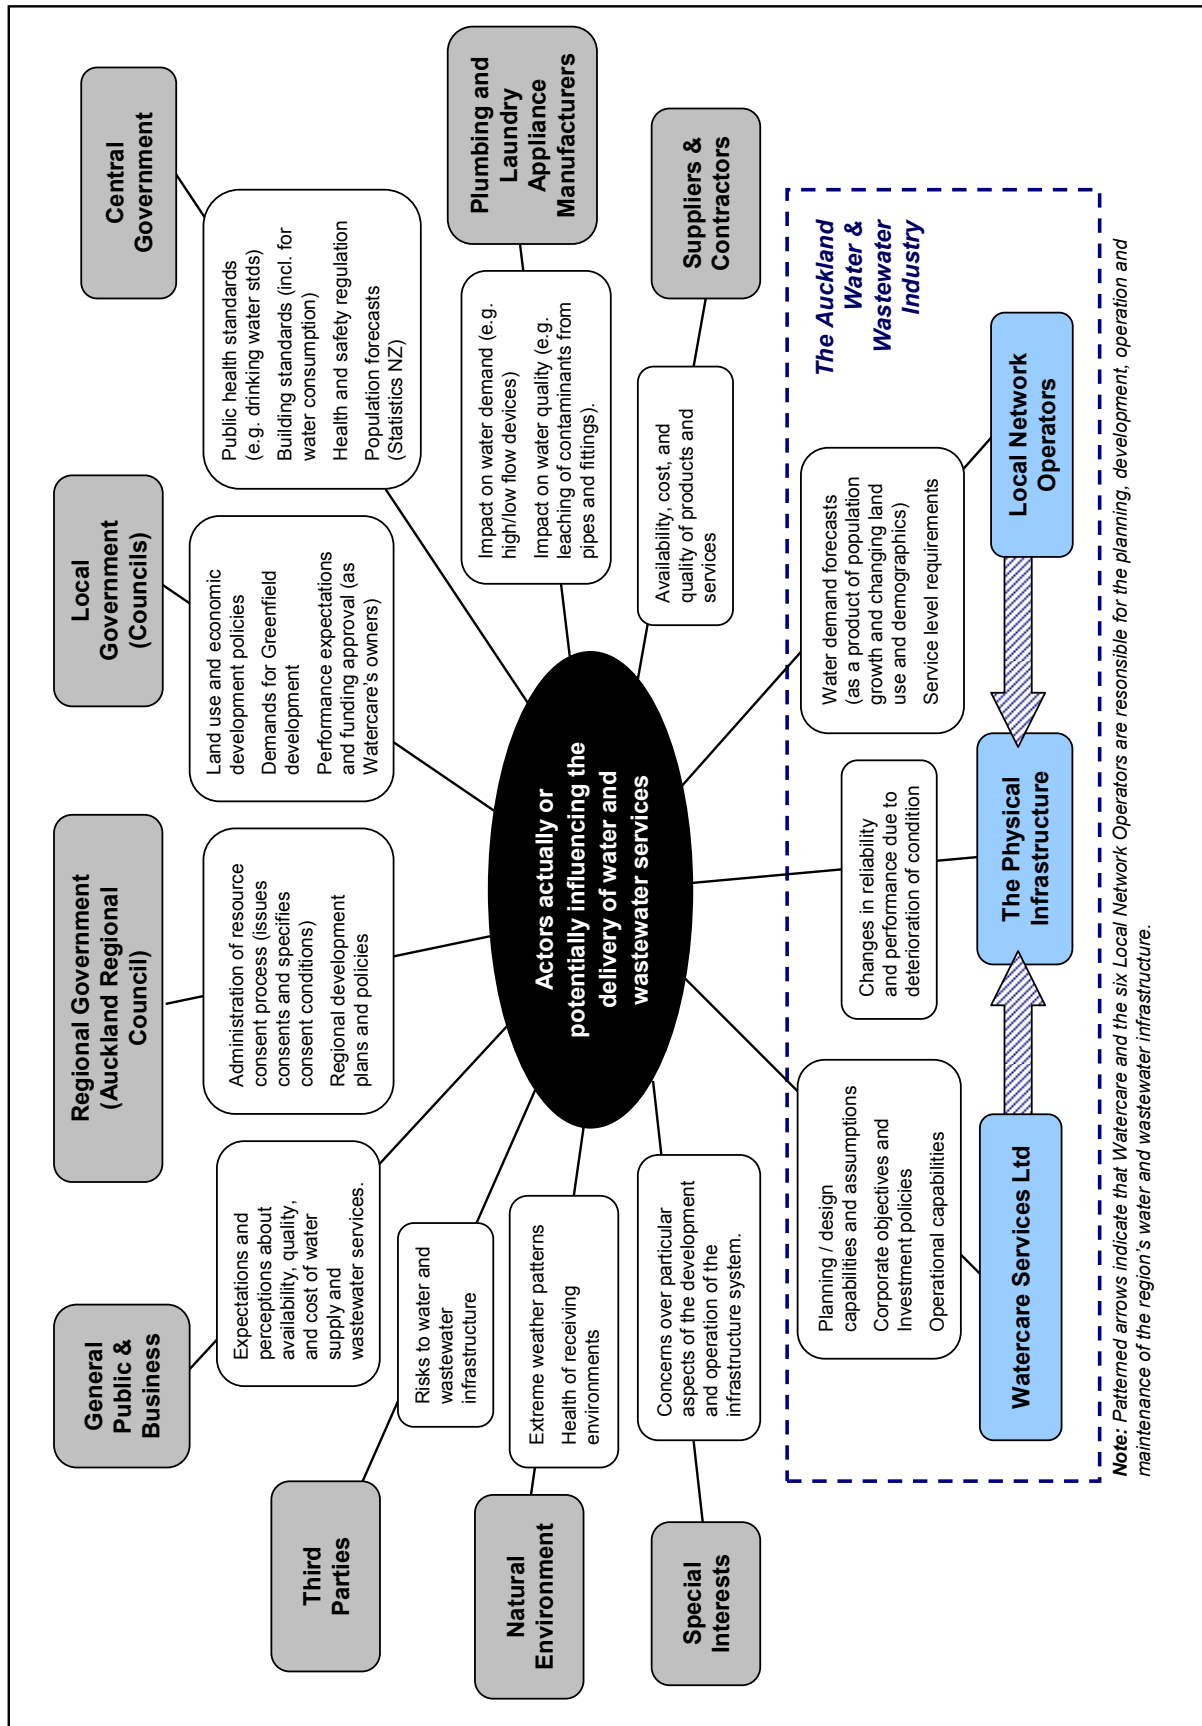


Figure II.4. Indicative mapping of actors and forces actually or potentially influencing the development and delivery of water and wastewater services in the Auckland region.

The other important feature was that the technological and organisational structures of the system were relatively static. At the time this research was conducted the last significant change to the organisational structure of the industry had occurred during the local government reforms of the early 1990s, when Watercare was created and vested with the water and wastewater assets of the old Auckland Regional Authority. The present structure at the time (Figure I.1, p 344) could not be altered without a legislative amendment<sup>24</sup>.

The technological structure of the system was also very slow to change. Once an asset or system was constructed or implemented, it was operated largely unchanged for the duration of its serviceable life, the length of which varied greatly depending on the type of asset: for example, dams and raw water tunnels > 100 years, water mains and sewer interceptors ~ 30 – 70 years, pumps ~ 10 – 20 years, electrical equipment < 10 years, computer and IT equipment < 5 years. Thus different parts of the infrastructure network were renewed at different rates (for instance, some of the infrastructure components within Watercare's networks were over 100 years old). The important point is that while the performance of an asset or system could be improved to some extent over its lifetime through operational changes or upgrades, those improvements were fundamentally limited by the state of technology at the time the asset or system was constructed. Opportunities for major technological change only arose at the end of the serviceable life of the asset. In contrast to the descriptions of LTS evolution in the literature, in the local context, technological evolution occurred as discrete, localised changes within the network, none of which fundamentally altered the overall structure of the system. This made the technological form of the infrastructure strongly path dependent, which was clearly evident in the form and geographical extent of Auckland's water and wastewater networks, reflecting the history of the systems dating back to the construction of the first dams in the Waitakere Ranges in the early 1900s and the opening of the Mangere wastewater treatment plant in 1960 (Murdoch 1993; Watercare Services Ltd. 2003b).

This did not mean, however, that significant technological change did not occur, as was illustrated by the construction of the 38km Waikato River pipeline (completed in 2002) and the upgrade of the Mangere Wastewater Treatment Plant (completed 2003); see Appendix I. The Waikato and Manukau projects both produced significant step-wise changes in technology and the performance of key parts of the infrastructure networks, and both were the products of formal, top-down planning processes, involving extensive consultation with

---

<sup>24</sup> As a result of a Royal Commission review of the Auckland Water Industry, Watercare and the Local Network Operators were amalgamated in 2009 into a single regional entity, under the name Watercare Services Ltd.

stakeholders and the communities of the Auckland region. The projects also shared another common feature. They were driven by regional concerns where significant performance objectives for the system at large were altered. Under business-as-usual decision making neither project would have been justifiable. But in each case a major performance objective was changed, necessitating a search for new solutions; for Waikato, it was the raising of the regional drought security standard, and for Manukau it was the de facto raising, via political interest, of acceptable standards for the quality of wastewater discharges to the Manukau harbour.

Such cases of significant political impetus for change was not, however, a frequent feature of the Auckland water industry. Rather, the majority of the time technological change was driven through the normal business planning functions of Watercare and the LNOs. This is framed in the following section through the lens of a general systems management model for the Watercare organisation.

### **General systems management model of the Watercare organisation**

Figure II.5 (page 401) presents a generic systems management model of organisations, adapted and simplified from Beer's Viable Systems Model (Beer 1972, 1985; Espejo and Harnden 1989). The model reflects a conceptualisation of the organisation as constituted by a variety of hierarchical, recursive, interacting systems, combining both human actors and heterogeneous physical components.

The first systems level consists of those systems which directly generate the primary productive outputs of the organisation. Since Watercare's core business was capturing, treating and delivering drinking water to the city, and collecting, transporting, treating, and disposing of the city's wastewater, the primary productive processes were the mechanical, electrical, chemical, and biological processes at work within the physical infrastructure system, as well as the daily operational activities performed by human actors.

The second systems level consists of those systems that are primarily concerned with the management (i.e. functional organisation) of the primary productive processes of the organisation. At Watercare these included maintenance planning and activities, asset performance monitoring and condition assessment, as well as asset planning and project development.

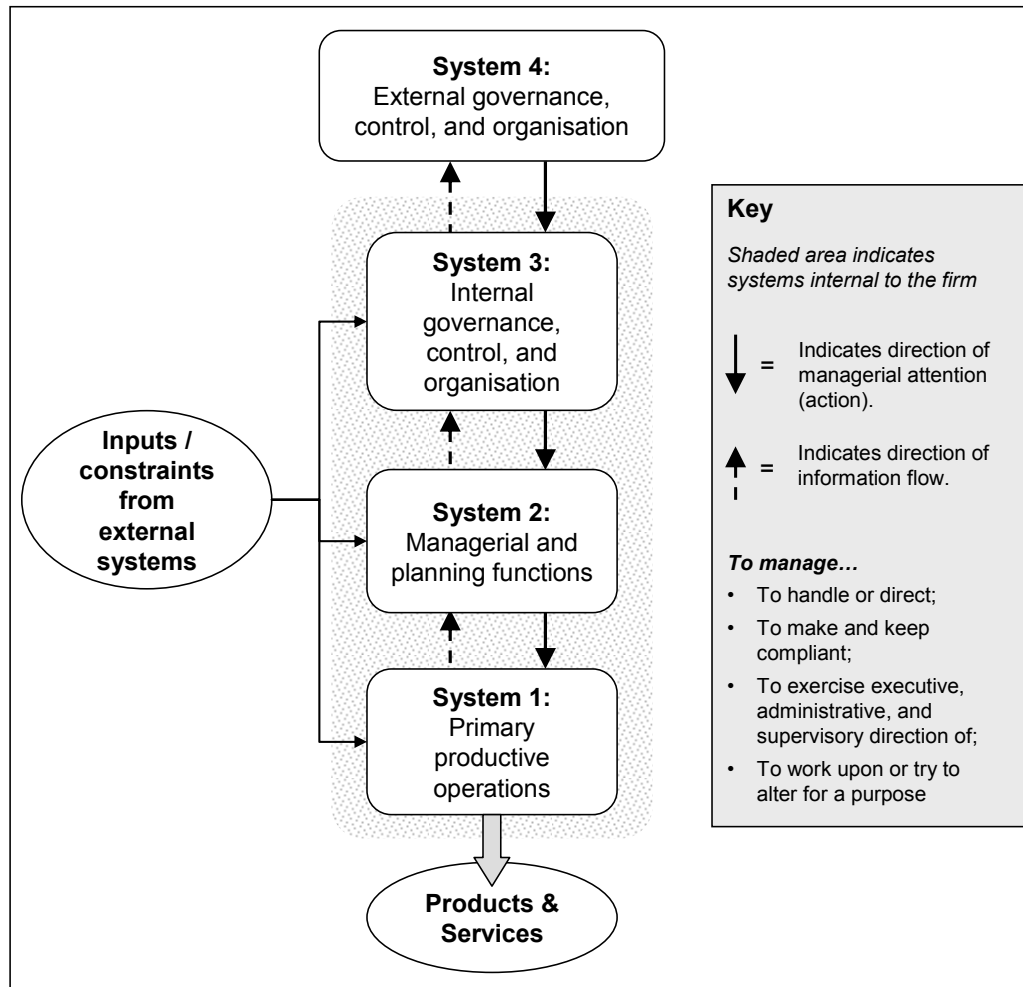
The third and fourth systems levels consist of those systems that are primarily concerned with ensuring that the organisation reliably delivers the required outputs to the required performance levels over time, i.e. with governance and control. Internally (systems level 3)

these would include internal audit, control, and oversight, policy making, and organisational restructuring, while externally (systems level 4) these would include the company's owners and the institutional mechanisms for the governance of Watercare's Board and executive management, and the processes by which Watercare engages with key stakeholders (e.g. customers, interest groups etc.). In this sense, the focus of the third systems level is with the internal structure, relationships, and performance of the various functional units that make up the company, while the focus of the fourth systems level is with the relationship between the company and its stakeholders, as well as the structure and institutional relationships of the wider industry of which the company is a part.

Figure II.5 thus represents a hierarchy of management control. Higher level systems define the objectives for, and organise and manage lower level systems (or parts thereof; this is represented by the solid arrows in Figure II.5). In order to do this the higher level systems must acquire information about the lower level systems, about their performance, and about their relationships with other systems, including the external environment (this is represented by the dotted arrows in Figure II.5).

In this model, changes in the functional arrangement and performance of each system are a product of managerial attention from the levels above. Two aspects of Watercare's operating context were relevant in this regard. First, the nature of the physical water and wastewater infrastructures was such that opportunities for major technological change were rarely, if ever found at the operational level. Further, when major technological change did occur, it was always localised. Second, Watercare's overarching mission was static (defined by legislation), as was the industry governance structure (Figure I.1, Appendix I, pg 347). The specific performance objectives of the system components were, however, subject to periodic, if infrequent, revision. As was noted in Appendix I, the dominant drivers of Watercare's investment programme, at a regional level, were the underlying trends of urban development, land-use change, and population growth across the region, regulatory changes affecting the company's levels of service, and the idiosyncratic needs and demands of the LNOs. These drivers occasionally led to significant technological changes, as illustrated by the Waikato and Manukau projects.

In this context, the primary mechanisms of technological evolution in the system were the formalised asset planning processes performed individually and collaboratively by Watercare and the LNOs (i.e. systems level 2). With respect to my research interests, this suggested that I should look for opportunities to influence those processes within the Watercare organisation.



*Figure II.5. Generic management systems model of organisations, after Beer's Viable Systems Model (Beer 1972, 1985; Espejo and Harnden 1989)*

## The nature of engineering planning and design

In abstract terms, a system may be defined as a group of components (which may be systems themselves) connected in a configuration that allows the system to perform specific functions within a defined boundary across which inputs and outputs flow (Ayyub 2003, Chapter 3). The system relies on inputs of materials and energy from its environment in order to function, and produces outputs as a result of that functioning. One may distinguish between a “system”, as a group of components and their relationships (rules of interaction), and the processes of which that system is a part (Emblemsvåg and Bras 2000). The term “process” means “a series of actions or operations conducing to an end”, and refers to the actual functioning of a system, including all of the actions and interactions of the system components, the inputs and their transformations, and the outputs produced. In this sense, “system” is taken to be a static concept with no temporal component in contrast to “process”

as a dynamic concept referring to the behaviour of systems. Processes transcend the boundaries of systems (i.e. a system may perform functions within several processes) such that, from a process perspective, the boundaries of a system are meaningless unless they also constitute the boundaries of the process(es) of which that system is a part (Emblemsvåg and Bras 2000). However, since it is function (action) that relates components in a system, a strict distinction between “system” and “process” is difficult to maintain. A system without function is simply a group of objects which exist, and is therefore not a system at all. For the purposes here, the terms “system” and “process” will be used relatively interchangeably, since systems imply processes and vice versa.

Scientific inquiry seeks to answer the question of “What is?” with regard to physical, natural, and social systems and processes, i.e. to investigate and produce knowledge of phenomena in the world so as to explain nature and society as they really are. Engineering design, in contrast, is practical (functional) in its intent. It seeks to answer the questions of “What for?”, “How to?”, and “How good?” in order to define the purposes, requisite functions, and performance criteria of the artefact or system being designed, and the form of, and means to produce that artefact or system (Auyang 2004, pp 153-155). Seeking answers to these questions constitutes an inquiry for framing and solving real-world problems.

### **Problem solving in engineering: the science and art of design**

Problems begin in the perceived undesirability of the current state of affairs relative to some set of requirements for how things ought to be; although these things need not be explicit to begin with – the problem solving process is initiated with the generation of doubts and the posing of questions, a general becoming awareness that something isn’t right (Buenaño 1999). Engineering design embodies a structured, systematic approach to problem framing and solving, usually characterised as a step-wise process involving iterative loops of analysis and synthesis. From the initial recognition that a problem may exist, the process proceeds through a number of stages (Pahl et al. 2007):

- i) Systems analysis to establish the nature (context, causes, antecedents, and consequences) of the problem;
- ii) goal formulation to define the objectives to be achieved, the relevant boundary conditions and constraints, and the criteria by which success will be evaluated;
- iii) solution development and evaluation where solution variants (systems, products, projects) are formulated, analysed (modelled), and evaluated to determine which



solution offers the best course of action for achieving the stated objectives; and  
iv) decision and implementation.

Engineering design is predicated on a functionalist systems methodology (Jackson 2000). The key aspects of this methodology are (Auyang 2004; Ayyub 2003; Jackson 2000; Pahl et al. 2007; Sydenham 2004):

- An assumption that the world is systematic and that real-world systems may be understood by examining the regularities in the relationships between sub-systems and the whole;
- An assumption that the analyses of problems may be conducted in systems terms, including the construction of models aiming to capture and represent real-world systems; and
- An assumption that systematic, rational approaches to problem structuring and analysis are more successful at producing effective solutions to problems than are ad hoc approaches which rely primarily on intuition and chance.

The systems framework plays an important role in formulating inquiry in engineering design by providing a methodology for ordering both our knowledge of the world, and our actions in it:

A generalized systems formulation allows scientists and engineers to develop a complete and comprehensive understanding of the nature of a problem and the underlying physics, processes, and activities. . . System definition is usually the first step in an overall methodology formulated for achieving a set of objectives. This definition can be based on observations at different system levels that are established based on these objectives. The observations can be about the different elements (or components) of the system, interactions among these elements, and the expected behavior of the system. Each level of knowledge that is obtained about an engineering problem defines a system to represent the project or the problem. As additional levels of knowledge are added to previous ones, higher epistemological levels of system definition and description are attained which, taken together, form a hierarchy of system descriptions. (*excerpt from Ayyub 2003, p 8*)

However, within the broader systems framework, engineering design is not simply the rational, mechanical application of theoretical knowledges. Rather the process proceeds as an iterative back-and-forth between moments of creativity and moments of objectivity, which reflects the complex nature of solving design problems (Auyang 2004; Pahl et al. 2007).

For any given problem the problem frame is not something that the problem itself possesses, but is a function of the relationship between the problem and the problem solver

(Buenaño 1999; see also Chapter 6). The ability to frame (and hence solve) a problem depends on the resources at hand, specifically the available knowledge of the problem domain, the available problem solving techniques, and the basic cognitive abilities of the problem solver, such that it is the behaviour of the problem solver which indicates whether the problem has been successfully framed – a problem is framed when the problem solver is readily able to identify a promising solution strategy (Buenaño 1999). In this sense, problem solving is a personal endeavour. Two people may approach the same problem in quite different ways and be more or less successful at it depending on their intelligence, the extent of their factual and experiential knowledges, and the sophistication of the problem solving resources they can bring to bear. But this does not mean that problem solving is an individual undertaking. On the contrary, problem solving is frequently a collective activity, and always involves the marshalling of calculative resources (Callon 1998b). But understanding (knowing) the problem is always personal (Buenaño 1999; Polanyi 1962). The corollary is that some problems may resist framing by all problem solvers due to limitations in the state of knowledge and resources, or may become subject to multiple, conflicting frames due to the plurality of perspectives that different problem solvers apply (Rittel and Webber 1973).

Each of the stages in the engineering design process therefore involves both the production and creative application of certain knowledges (factual and experiential) in a particular domain of interest. In the analysis phases complex systems are investigated and resolved into their constituent components such that the components, and their behaviours and interrelationships may be studied and understood (Pahl et al. 2007). Analysis calls for identification, definition, structuring, and arrangement of system elements, i.e. modelling. Synthesis, in contrast, is where engineers call on their acquired factual and experiential knowledges to compose new arrangements which ought to produce certain desired effects, i.e. to propose possible solutions (Pahl et al. 2007). Since these arrangements are propositional, synthesis is followed by further analysis to evaluate how effective the designs ought to be in achieving the overall objectives for the system. In this way, engineering designs are developed via iterative phases of analysis and synthesis until an optimal solution is attained within the constraints of the problem domain (Pahl et al. 2007).

Einstein commented, with regard to the production of scientific theory, that the conception of a theoretical model from a body of experiences was not a logical deductive development but rather an intuitive leap – the theoretical idea “does not arise apart from and independent of experience; nor can it be derived from experience by a purely logical procedure. It is produced by a creative act” (Einstein 1954, quoted in Auyang 2004, p 158). The same is true in engineering design. The imagining of novel solutions or approaches to

engineering problems involves insight and creativity (Auyang 2004; Hurst 1999; Pahl et al. 2007). It is not a purely logical progression from the knowledge at hand to the form of a solution or strategy. However, once that form has been imagined, it must be made objective as calculated plans. Here the design process switches back to analysis, where the necessary analytical tools are brought to bear on the design concept so as to sharpen it up and make it concrete, to make it specific and precise (quantitative), and to demonstrate that it will work; a movement seen quite clearly in the progression from concept design through to detailed design in engineering projects (Auyang 2004; Hurst 1999; Pahl et al. 2007; Sydenham 2004). In this endeavour of making concrete, design ideas are developed from general concepts (e.g. a pipeline from point  $X$  to point  $Y$ , with a capacity of  $Z$  m<sup>3</sup> per second) to detailed plans of a precise and quantitative nature (e.g. a pipeline from point  $X$  to point  $Y$ , following a particular specified route, constructed from material  $M$  with lining  $L$ , varying in diameter from  $D_i$  to  $D_n$  over the route length... and so on).

In summary, engineering design is a creative yet structured problem solving process, the outcome of which is technology to meet human needs. Conceptualised as an iterative, step-wise progression from the conceptual (qualitative) to the concrete (quantitative), it employs a functionalist (i.e. means-ends) systems methodology, embodying a continual back-and-forth between analysis and synthesis, creativity and objectivity, between art and science (Armstrong 2008; Auyang 2004; Frost 1992; Hurst 1999; Jackson 2000; Pahl et al. 2007; Sydenham 2004).

## Identifying potential contexts for inquiry

Watercare's capital planning and project development processes were the fundamental processes through which the water and wastewater infrastructure evolved. In this regard, those processes may be considered important dynamic capabilities; a dynamic capability being the capacity of an organisation to purposefully extend or modify its resource base (Helfat et al. 2007). The concept of dynamic capability includes the capacity with which to identify the need or opportunity for change, formulate a response to such a need or opportunity, and implement a course of action (Helfat et al. 2007). Key features of the notion of dynamic capability include: (i) a repeatable, patterned, and somewhat practiced activity, (ii) a purposeful activity, i.e. one that is carried out with specific intent, and (iii) an activity that creates new resources or extends or modifies existing resources, where 'resources' may be understood as anything that the organisation draws upon to achieve its aims (Helfat et al. 2007, pp 2 – 7). Although dynamic capabilities are inherently innovative, these features

distinguish dynamic capabilities from change which occurs as a result of ad hoc or idiosyncratic activities. This does not deny that experimentation and accidental discovery are important sources of innovative change, but rather recognises that effective change requires strategy and management (Helfat et al. 2007).

By modifying and extending the resource base of an organisation, dynamic capabilities have the potential to create value and competitive advantage, and are therefore understood as one factor contributing to the evolutionary fitness of organisations in competitive environments; although whether a dynamic capability creates significant value, and whether that value translates into sustainable competitive advantage is heavily context dependent (Helfat et al. 2007, Chapter 1). Watercare and the LNOs operated in a non-competitive, public sector arena where the term “value” had a different meaning than in normal market contexts (i.e. value-for-money; The Treasury 2008). Nevertheless, the underlying relationship remains the same – dynamic capabilities contribute to long-term performance.

The benefits that dynamic capabilities confer in terms of sustaining and improving organisational performance depend on the efficacy of the various managerial and organisational processes which are constitutive of dynamic capabilities. That is, the identification of a need or opportunity “involves problemistic search and opportunity recognition processes”, formulation of a response “involves internal selection processes and resource allocation processes”, and implementation “involves a variety of managerial and organizational processes, depending on the nature of the objective and the specific tasks required” (Helfat et al. 2007, Chapter 3 with C. A. Maritan, pp 30-31). The development and improvement of dynamic capabilities therefore requires a context-specific understanding of these underlying processes.

As was described earlier, Watercare’s capital planning occurs at two levels: (1) master planning, which describes the strategic development of the infrastructure system over the long term, and (2) project planning and development, which details specific changes to be made to the system in the short term. Master planning evaluates the need for and timing of individual projects within the broader context of the development of the system as a whole, while project planning produces detailed design plans for individual capital investments. Three aspects of Watercare’s capital planning capabilities stood out as potentially fruitful contexts in which to develop an inquiry.

The first was what might loosely be called the strategic planning framework. In simple terms, the strategic development of the system is the product of a supply – demand equation, where “demand” encompasses anticipated future production, changes to levels of service and regulatory conditions, and anticipated greenfield developments in the city, and “supply”

encompasses the capacity and reliability of the existing infrastructure (which change over time as assets deteriorate), the evolving state of technology, and the resources available to construct new infrastructure, and renew or replace existing infrastructure. The strategic planning framework encompasses the various objectives, inputs, assumptions, and methods used to formulate this “equation”, i.e. to conceptualise and calculate what the infrastructure system might look like in the future. In this sense the framework is not a fixed entity, but rather evolves over time. It might reasonably be expected that for any given performance of the strategic planning cycle, different combinations of objectives, inputs, assumptions, and methods would produce different plans for the long-term development of the system (although the possibilities for variation would clearly be limited in practice). Accordingly, the first option for inquiry was to investigate the nature of Watercare’s strategic planning framework, or some aspect of it, with a view to understanding where and how to intervene to reorient the framework to address new outcomes. One possibility, for example, was to investigate how scenario planning could be used to challenge objectives and assumptions.

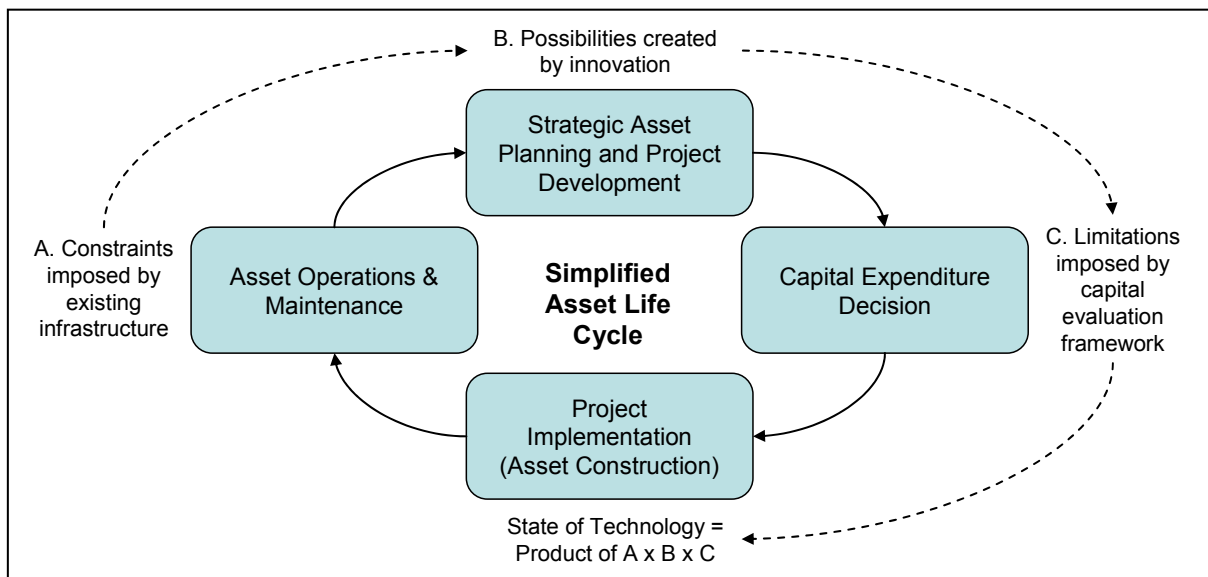
The second relevant context was the definition phase of the project planning process. The development of a project from concept through to detailed plans for construction is fundamentally a process of engineering design. Although the process may be conveniently represented as a rational, step-wise procedure culminating in a discrete decision event (e.g. Figure I.3, pg 355), in practice the processes preceding formal sign-off of a CapEx are complex and collaborative, and characterised by social interaction, politics, and discursive negotiation between actors, where technical engineering knowledge is represented alongside competing cultural, economic, and political knowledges. A particular feature of engineering design is an iterative back-and-forth between moments of creative synthesis and moments of objective analysis (i.e. between having an idea, and figuring out whether that idea will actually work in practice). Although these creativity-analysis loops occur throughout the design process, creativity is more prominent at the beginning when options are being imagined, while analysis is more prominent toward the end when details are being worked out. Accordingly, important possibilities for facilitating innovation might be found in the contingency and experimental nature of problem solving which characterises the micro-scale practices of engineering design. By challenging conventional thinking at the beginning of the planning process, before designs became ‘locked in’, innovative ideas could be encouraged and brought to the fore. Although conventional thinking would probably still win out in most cases, the expectation is that if the company culture encourages ‘out of the box’ thinking on a systematic basis, then there is a higher likelihood that innovative designs will make it to implementation. The second possibility for inquiry was thus to investigate the early phases of

project development, with a view to understanding the factors influencing innovation in design, and thus where and how to intervene to encourage innovation in project planning.

The third relevant context, closely related to the first above, was what might loosely be called the capital decision making framework. Once project options have been identified they must be evaluated against decision criteria in order to identify the rational choice. At any point in time the decision making rationality of the firm is made at least partially explicit in certain policies and frameworks, which embed prior decisions about objectives and their relative priorities, and about the accepted means of calculating and representing investment outcomes. At Watercare, for example, these included the Net Present Value spreadsheets and the company's risk framework, along with a formal framework for mapping capital expenditure against three fundamental drivers of the business: growth, renewal, and service level improvement. Of these, the risk framework was arguably the most important because it defined the relative priority of both capital projects and maintenance activities across the company, i.e. the higher the risk score, the higher the priority, and it oriented attention to certain performance outcomes, i.e. if outcomes could not be evaluated within the framework then they were de facto not important and invisible for the purposes of decision making. The Waikato and Manukau projects, described earlier, were clearly extreme cases where major technological change became justifiable in light of significant political willingness to create a new status quo. Nevertheless, they suggested that important possibilities for facilitating change might be created by opening up reflection and discussion about decision making infrastructures, and particularly about performance objectives and their relative priorities (recognising that although cost and risk will remain central criteria, Watercare's corporate objectives cover a much broader range of outcomes). Accordingly, the third opportunity for inquiry was to investigate the decision making rationality of the firm, with a view to understanding how business cases for investment are justified.

The basic logic was that the three contexts identified above represented important sources of innovation and control with respect to Watercare's primary dynamic capability – the capital planning process. In essence, the implementation of projects (construction of assets) can produce stepwise changes in the performance of the system (i.e. the replacement of old technology with new). The nature and magnitude of those changes is a function of the state of technology selected during the planning process. This is, in turn, a function of: (A) the constraints imposed by the existing infrastructure, (B) innovation during the planning process (at both strategic and project levels), and (C) the criteria by which capital expenditure is justified. The latter criteria are, in turn, a function of the corporate objectives and strategy. If the objective of the research were to influence the performance of the

company through the capital planning process, then (B) and (C) represented the important contexts (above) in which to investigate opportunities for doing so (see Figure II.6).



*Figure II.6. Possibilities and limitations on the state of technology*

## Recognising the opportunity presented by the CRM's project

In early 2007 I became aware of the new Corporate Risk Manager's intentions to redevelop Watercare's Risk Management functions, and realised that I could formulate an inquiry around that project. But the circumstances meant that I had little time to properly define the focus, scope, and parameters of the inquiry before embarking on the empirical data collection. As I initially conceived the project, the focus of the research was not actually the role of the Corporate Risk Manager per se. Rather, the research objective was to describe and explain how "risk", as an important decision making concept, and the frameworks, processes, and practices associated with it, changed over time. I framed the study as an exploration of how the framework of "risk" evolved in the organisation. I posed a series of open-ended questions that were intended to focus my attention on how the concept of risk was understood and perceived by actors within the company, how it was constructed through the tools, practices, and discourses of risk management, what roles it played in capital decision making, how concepts of risk and risk management practices had evolved historically, and how they would change as a result of the CRM's redevelopment programme. In this sense, I saw the CRM as simply one actor among many, though perhaps one with more influence than most over the development of risk management in the company.

In August 2007 I identified that the task of implementing ERM and the CRM's experiences with performing that task should become the primary foci of the inquiry. This re-orientation was the product of two factors. First, I identified that the implementation of ERM and the role of Chief Risk Officer had garnered little attention from academics, and that I was in a uniquely privileged position to longitudinally follow, describe, and analyse the performance of a CRO engaged in the implementation of ERM. Second, because I was reading across a range of literatures, I had uncovered a diversity of perspectives on risk and risk management, but no single perspective provided sufficient understanding to make sense of everything that seemed to be relevant in the domain in which I was located. Each of those literatures offered a different way of seeing the world and understanding what was going on when the CRM intervened in the Watercare organisation. I realised that these different perspectives needed to be integrated around the role of the Chief Risk Officer in order to inform the task of implementing ERM. Although I would remain uncertain for quite some time about what specific contributions the thesis might make, by the end of August 2007 I had identified, at least very generally, both practical and academic problems to address, and theory that was relevant to the practice context. The CRM's performance therefore became the primary object of my attention (the methodology of the inquiry is discussed in detail in Chapter 2).



# Appendix III

## Review of literature around ERM and the CRO role

### **Drivers of ERM uptake**

Two different but convergent pressures may be discerned behind the emergence and uptake of enterprise risk management (Dickinson 2001; Power 2005a, 2007). One is responsibility-based, rooted in the corporate governance revolution mentioned above (Power 2005a, 2007). This motivation posits risk management as a technology for the good governance and internal control of the organisation (Julien and Rieger 2003; Lam 2003; Miccolis and Shah 2000; Ward 2006). The other is value-based, where risk management is seen not just as a compliance function protecting shareholder value, but also as a positive force for creating it (Barton et al. 2002; KPMG International 2007; Lam 2003; Meulbroek 2002; Nocco and Stulz 2006).

### **Good corporate governance**

The origins of the corporate governance revolution can be traced to the mid 1980s. Banking failures were already a global phenomenon (Basel Committee on Banking Supervision 2004), and corporate governance and internal control were becoming matters of regulatory concern for the financial industry. The late 1980s and early 1990s saw the publication of the Report of

the National Commission on Fraudulent Financial Reporting (and accompanying guidelines) in the U.S. (NCFFR 1987)<sup>25</sup>, and of the Report of the Committee on the Financial Aspects of Corporate Governance (and accompanying code of best practice, the ‘Cadbury Code’) in the U.K. (CFACG 1992)<sup>26</sup>. The studies represented in these reports were motivated by concerns, in both countries, about the quality of corporate governance, accounting, and auditing practices, and the lack of accountability at board level; concerns which had been heightened by recurring corporate scandals and failures (CFACG 1992; NCFFR 1987). Both reports were the products of private sector initiatives, but both had regulatory impact. The Treadway Commission appeared twice before U.S. Congressional Hearings on the adequacy of auditing, accounting, and financial reporting practices under federal securities laws, and the London Stock Exchange required companies to issue a statement of compliance with the Cadbury Code (including reasons for areas of non-compliance) as part of continuing listing obligations (CFACG 1992; NCFFR 1987). The NCFFR and CFACG studies were the first in the corporate governance and risk management revolution which took place in the 1990s and continued into the new millenium. They were followed by a raft of studies, reports, frameworks, and legislative initiatives seeking to provide guidelines for and improve oversight of corporate governance, internal control, and risk management in both private and public sectors around the world (see Table III.1).

Power (2007) argued that the governance explosion focused and intensified policy and regulatory attention on the internal organisation and design of large entities, thereby centering and emphasising the importance of internal control. He notes that “[w]hile the role of external auditing, monitoring and inspection has not been superceded, demonstrable capacities for self-control and self-observation have grown in regulatory importance” (Power 2007, p 60). As the governance revolution unfolded, ideas about internal control and risk management became increasingly fused together, into what Power (2005a, 2007) has called the emergence of a control-based concept of risk management. The COSO framework of 1991, for example, explicitly positioned risk management as a pre-condition for internal control, “Control activities are the policies and procedures that help ensure... that necessary actions are taken to address risks to achievement of the entity’s objectives.” (COSO 1991). The

---

<sup>25</sup> Also known as the Treadway Commission, sponsored by the The Committee of Sponsoring Organisations (COSO), which was established in 1985, and consisted of five major professional associations in the United States, the American Accounting Association, the American Institute of Certified Public Accountants, Financial Executives International, The Institute of Internal Auditors, and the National Association of Accountants (now the Institute of Management Accountants). See <http://www.coso.org>.

<sup>26</sup> Sponsored by the Financial Reporting Council, the London Stock Exchange and the accountancy profession.

explicit linking of notions of internal control and risk management to the achievement of an organisation's objectives and the protection of shareholder value effectively repositioned them as "core values... for corporate governance" (Power 2007, p 60). More recently, internal control has been subsumed as just one functional area within the broader scope of Enterprise Risk Management (see, for example, COSO 2004), which reflects the continuing evolution of ideas about corporate governance: "Corporate governance is evolving from command-and-control dictums to a continuous process that assesses, sources, measures and manages risks across the enterprise" (Julien and Rieger 2003, p 34).

*Table III.1. Selected corporate governance, control, and risk management initiatives*

<b>Year</b>	<b>Country &amp; Initiative</b>
1987	<b>US:</b> COSO, Treadway Commission: <i>Report of the National Commission on Fraudulent Financial Reporting</i>
1991	<b>US:</b> COSO: <i>Internal control – Integrated Framework</i>
1992	<b>UK:</b> The Cadbury report: <i>Report of the Committee on the Financial Aspects of Corporate Governance</i>
1993	<b>US:</b> Group of Thirty Report <i>Derivatives: Practices and Principles</i>
1994	<b>Canada:</b> The Toronto Stock Exchange Dey Report: <i>Where Were The Directors? Guidelines for Improved Corporate Governance in Canada</i>
1995	<b>Australia / New Zealand:</b> <i>AS/NZS 4360:1995 Risk Management</i> (updated 1999 and 2004)
1995	<b>Canada:</b> Canadian Institute of Chartered Accountants Criteria of Control (CoCo) Board: <i>Guidance on Control</i>
1998	<b>UK:</b> The Hampel report: <i>Committee on Corporate Governance: Final Report</i>
1998	<b>UK:</b> Financial Services Authority: <i>The Combined Code</i> (revised 2003)
1998	<b>Germany:</b> <i>The Control and Transparency Act (KonTrAG)</i>
1999	<b>UK:</b> Institute of Chartered Accountants in England and Wales, the Turnbull Guidance: <i>Internal Control: Guidance for Directors on the Combined Code</i>
2002	<b>US:</b> <i>Public Company Accounting Reform and Investor Protection Act</i> (Sarbanes-Oxley)
2003	<b>Australia:</b> Australian Stock Exchange's Corporate Governance Council: <i>Principles of Good Corporate Governance and Best Practice Recommendations</i>
2004	<b>US:</b> COSO: <i>Enterprise Risk Management – Integrated Framework</i>
2004	<b>Global:</b> <i>International Convergence of Capital Measurement and Capital Standards</i> (Basel II)
2004	<b>Global:</b> <i>OCED Principles of Corporate Governance</i>

Sources for Table III.1: (Aabo et al. 2005; Lam 2006; Manifest 2004; Miccolis and Shah 2000; Power 2007; Tarantino 2006; Vinten 2001)

Although Power (2005a, 2007) suggests that the popular claims for the functionality of ERM are difficult to realise in practice, ERM has attracted considerable attention and commitment from companies around the world, in both private and public sectors, as a response to global pressures for good governance and control. Industry surveys on risk management (reviewed later) found that regulatory compliance pressures, particularly around financial reporting and corporate governance, and increased expectations from shareholders and stakeholders for greater assurance, accountability, and better risk management, were the most significant drivers of corporate interest in developing ERM capabilities (Brancato et al. 2006; Ernst & Young 2006a, 2006b; KPMG International 2007; Miccolis, Hively, and Merkley 2001; Towers Perrin 2006). As a result, Boards and senior executives are now perceived to have greater awareness of their accountabilities and responsibilities for risk, and are more involved in risk oversight (Brancato et al. 2006; Ernst & Young 2006b).

### **Shareholder value creation**

The other driver behind ERM uptake is the proposition that, beyond preserving value and meeting governance and compliance demands, enterprise risk management can actually contribute to the creation of competitive advantage and shareholder value. For example:

Leveraging risk management to make better business decisions by incorporating risk/return considerations in product development and pricing, relationship management, investment and portfolio management, and mergers & acquisitions... can be a powerful tool for improving business performance... and thus maximise shareholder value. (*excerpt from Lam 2000, p 4*)

The implementation of ERM programmes by organisations around the world is increasingly being driven by such recognition (or perception) of the value-add potential of ERM. In Barton et al.'s study of ERM practices in five large U.S. corporates, they note: "One common theme emerged. Each company believed it was creating, protecting, and enhancing value by managing enterprise risk" (2002, p 11). The Towers Perrin survey reflects a similar finding, "[t]he value proposition of ERM was very clearly agreed upon among executives. Above all, the key drivers of ERM adoption include the economic value it generates and the improvement in a firm's risk-return relationships" (2006, p 6).

Since regulatory and other assurance requirements are growing, and the costs of compliance with those requirements are significant, companies are searching for ways to move beyond mere compliance and extract value from those investments (KPMG International 2007). KPMG International's Global Head of Internal Audit Services, Mike

Nolan, contends that, globally, the key questions being asked of corporate risk and controls functions are: “How can we transform an expensive compliance obligation into a real business advantage? How can we deliver significant and quantifiable operational and financial value from the risk spend? How do we reconcile increased efficiency with increased risk and controls management?” (KPMG International 2007, p 5). The argument, in essence, is that the company that secures even a minimal value-add from a compliance investment should have a competitive advantage over one that does not.

More than anything else, it is the improvements to strategic decision making that companies appear to value most highly. Companies, both locally and globally, perceive that they are facing increasing levels of risk in their operating environments. Although, as Ernst & Young (2006b) note, some of this perception may be due to reduced tolerances for risk, companies are perceiving increased threats from natural disasters (e.g. Hurricane Katrina) and potential pandemics (e.g. avian influenza), increased regulatory attention and liability risks, and new business risks arising from increased reliance on technology and complex infrastructures, greater supply chain complexity, and exposures from pursuing growth opportunities in ‘emerging’ markets (Ernst & Young 2006b; KPMG International 2007; Towers Perrin 2006). In this context, companies are turning to enterprise risk management to provide more effective decision making:

Once largely focused on avoiding loss and complying with regulations, risk and controls are now increasingly required to show that they also add value... [this] typically means that these functions participate fully in broad strategic issues, such as mergers and acquisitions, as well as contributing to more focused business decisions, such as those related to product development. Risk management, in short, is increasingly seen as a partner to the operating business as well as a key strategic tool of the Board. (*excerpt from KPMG International 2007, p 6*).

Within the ERM literature, there are a number of common claims about how ERM leads to (or may lead to) value creation. These include knowing how disparate risks from across the business relate to each other and whether they offset or compound each other, knowing how risks could affect both financial position and earnings, knowing how much capital is at risk and being able to account for this in resource allocation and product pricing decisions, aligning risk appetite and strategy, reducing earnings volatility and hence tax liabilities while simultaneously increasing attractiveness to investors, and proactively identifying and pursuing valuable investment opportunities (Barton et al. 2002; COSO 2004; Culp 2002; Lam 2003, 2006; Meulbroek 2002; Nocco and Stulz 2006). Respondents to industry surveys also emphasised the cost efficiency and effectiveness benefits perceived to derive from the explicit alignment of strategy with corporate objectives, and the cross-functional communication

and co-ordination that is achieved through ERM programmes (Ernst & Young 2006b; KPMG International 2007; Towers Perrin 2006).

In all of these value claims, the implication is that ERM improves the efficiency and effectiveness of decision making. Knowing more about the company's risk environment, its risk profile and appetite, and about what other parts of the company are doing and what risks they face should, in theory, lead to more informed decision making, while the formalisation and structure that ERM promotes should lead to more rigorous analysis and greater decision making discipline. All of which should, in turn, constitute a foundation for robust decision making and help generate greater confidence in the decisions reached, not only for the decision makers themselves, but also for those parties requiring assurance of the quality of corporate decision making.

This basic argument has been formally recognised by institutional investors and major ratings agencies such as Moody's, and Standard & Poor's. An Ernst & Young survey of 138 of the world's largest institutional investors found that 82% of those surveyed were consistently willing to pay a premium on share price for companies that demonstrate effective risk management practices (Ernst & Young 2007). The rating agency Standard & Poor's announced in 2007 that they would incorporate risk management analysis into their credit ratings process as a means to benchmark the quality of management judgement; thus explicitly endorsing the perception that good risk management leads to better decision making, and, in turn, improvements in business performance. They note:

Management analysis is arguably the most qualitative and immeasurable among the many considerations of a Standard & Poor's Ratings Services' credit rating. A rating committee's opinion of management's overall capabilities, fidelity to and consistency of a sound strategy, and adaptability to changing circumstances are perhaps the strongest influences on the future direction of a credit rating. Nevertheless, the quality of management judgment is not as easily benchmarked by quantitative metrics in the way that ratios and models of cash flow adequacy, liquidity, earnings capacity, and leverage help shape our views of a company's financial profile.... We now propose to introduce Enterprise Risk Management (ERM) analysis into the corporate credit ratings process globally as a forward-looking, structured framework to evaluate management as a principal component in determining the overall business profile... we expect that deterioration or improvement in a company's ERM quality would potentially drive rating and outlook changes before the consequences are apparent in published financial results. Companies with superior ERM should have less volatility in earnings and cash flow, and will optimize the risk/return relationship. (*excerpt from Dreyer and Ingram 2007, p 2*)

## Corporate interest in ERM and the CRO role

Throughout the 1990s and 2000s there was increasing corporate and public sector interest in and commitment to ERM, in both financial and non-financial industries. Corporate interest in ERM, and emerging experiences with its implementation, have been the subject of a number of industry surveys in recent years. Table III.2 (p 420) presents a comparative summary of statistics on the uptake and development of ERM programmes in industry, drawn from seven such surveys carried out prior to the 2008-09 financial crisis.

While the surveys are a rich source of information, it must be noted that interpreting and comparing the results is fraught with difficulty. Aside from the fact that the surveys employ different methodologies, and ask different questions, the main factor limiting comparison is the absence of a common definition of ERM. Each of the surveys employs the term “Enterprise Risk Management”, but not all of them define it, and none of them define a development profile for ERM, i.e. what the characteristics of ERM are in practice, and what differentiates a partially developed ERM programme from a well-developed or fully developed programme (this latter point is somewhat surprising given the proliferation of ERM Capability Maturity Models; see Appendix V). Consequently, the basis on which participants responded to questions about the state of ERM within their firms was left open to their individual understanding and interpretation of the ERM concept in their respective industries, and their perceptions of how ERM in their firms aligned with that understanding. Most of the surveys did, however, examine, via more specific questions, various particular aspects of the state of ERM programmes. These included: what motivated firms to pursue ERM (KPMG International 2007; Towers Perrin 2006); the degree to which ERM was strategically focussed (Marsh and RIMS 2007, 2008); the degree of alignment between risk management and strategic objectives (Ernst & Young 2006b); which risks firms were most worried about (Marsh and RIMS 2006; Miccolis et al. 2001); the tools and procedures firms are using to implement ERM and the barriers and challenges they are facing (KPMG International 2007; Miccolis et al. 2001); or the perceived benefits and degree of satisfaction with ERM (Marsh and RIMS 2008; Miccolis et al. 2001).

Surveys published before the global financial crisis of 2008-2009 reported a number of trends which were indicative of continuing corporate interest in Enterprise Risk Management; in the sense that proponents of ERM saw it as fulfilling corporate needs with respect to each of the following trends:

- A number of surveys reported that executives were increasingly looking for risk

management to become more strategic (Deloitte and Touche LLP 2007; Ernst & Young 2006b; KPMG International 2007; Marsh and RIMS 2007, 2008). Driven, at least in part, by common perceptions of an increasingly complex and volatile risk landscape, executives felt that risk management had to become broader in its focus and more effectively address less traditional risks such as brand/reputation risk, human capital risk, technology/e-risk, and business continuity risk and crisis management (Deloitte and Touche LLP 2007; Marsh and RIMS 2007, 2008). There was also an increasing focus on emergent risks, and particularly supply chain, vendor, and geo-political risks arising from globalisation and outsourcing trends (Deloitte and Touche LLP 2007; KPMG International 2007). These were all examples of 'high-importance' risks that executives felt are not well handled by their firms' existing risk management capabilities (Deloitte and Touche LLP 2007; Marsh and RIMS 2007, 2008).

- Achieving greater penetration and integration of risk management into business processes, facilitating risk-aware cultures, measuring and aggregating risk (particularly non-traditional risks), and developing tools and systems to support risk management were reported as the most immediate and significant challenges that firms were facing in developing their risk management capabilities (Deloitte and Touche LLP 2007; Ernst & Young 2006b, 2008; KPMG International 2007).
- There was an increasing awareness of and focus on risk at Board level, as well as greater participation by corporate Board's in risk oversight (Brancato et al. 2006; Deloitte and Touche LLP 2007; Ernst & Young 2006a, 2006b, 2008; PWC 2008).

Table III.2 presents the one statistic that is common to nearly all of the surveys – the number of firms who have implemented ERM or are in the process of doing so. The surveys in Table III.2 were identified via a Google search on the Internet, or from other literature sources, and were selected for inclusion in Table III.2 on the basis of three criteria: (1) the scope of the survey was preferably global, (2) survey respondents represented a variety of industries, and (3) the survey reported, in some form, the statistic just mentioned. The exception is the Deloitte and Touche LLP survey, which, while global in scope, was limited to the financial services industry. This survey was included in Table III.2 for comparative purposes because the financial services industry is generally considered to be more advanced in its implementation of ERM than other industries. The variation in results reported by the seven surveys is such that it is difficult to interpret trends over time. Nevertheless, the results



in Table III.2 would seem to support the following statements:

- The substantial majority of firms surveyed are planning, developing, or have developed ERM capabilities. The percentage of firms not interested in ERM – approx 1/4 to 1/3 of firms surveyed – is fairly consistent across the surveys.
- The percentage of firms claiming to have a fully developed ERM programme in place is, for the most part, fairly low (less than 20%), which is probably indicative of the continuing challenges firms face in implementing ERM (KPMG International 2007). This number is also widely variable across the surveys, which is, perhaps, indicative of a lack of stable understanding about what mature ERM should look like, and the evolution of ERM practices over time. The percentage of firms claiming to have a fully developed ERM programme in place in the financial services industry survey is considerably higher at 35%.
- The later surveys (2007, 2008) report a much lower percentage of firms considering or planning ERM (~20%) than the earlier surveys (~40%), which may be indicative of a recent plateauing of interest in developing ERM capabilities (Marsh and RIMS 2008).

Table III.3 (p 421) presents a comparative summary of industry surveys showing the penetration of the role of CRO amongst global corporates. As with Table III.2, the surveys in Table III.3 were identified via a Google search on the Internet, or from other literature sources, and were selected for inclusion in Table III.3 on the basis of three criteria: (1) the scope of the survey was preferably global, (2) survey respondents represented a variety of industries (the latter three surveys in Table III.3 were limited only to the global finance or insurance industries and are included for comparison), and (3) the survey reported the percentage of firms claiming a CRO or equivalent position. Bearing in mind the aforementioned caveats with respect to interpreting findings across different surveys, the following points can be drawn from Table III.3: (1) around one quarter to one third of companies in non-financial industries have a CRO position, and (2) the penetration of the CRO role in the financial services industry is substantially higher than in non-financial industries, at around 60 – 80%. In combination Tables III.2 and III.3 appear to support a positive correlation between the adoption of ERM and the creation of the CRO position in firms, while also indicating that not all firms with or considering ERM necessarily have a CRO position.

**Table III.2. Comparative summary of surveys on ERM uptake**

Survey	ERM Trends & Emerging Practices	Towers Perrin A Changing Risk Landscape	Excellence in Risk Management III	Excellence in Risk Management IV	The Evolution of Risk and Controls	Excellence in Risk Management V	Global Risk Management Survey, 5 <sup>th</sup> Ed.
Reference	Miccolis, Hively, and Merkley (2001)	Towers Perrin (2006)	Marsh and RIMS (2006)	Marsh and RIMS (2007)	KPMG International (2007)	Marsh and RIMS (2008)	Deloitte & Touche LLP (2007)
Number of respondents <sup>1</sup>	> 130	75	866	501	435	n/a <sup>2</sup>	130
Scope <sup>3, 4</sup>	Global, multi-industry	U.S., multi-industry	U.S., multi-industry	U.S., multi-industry	Global, multi-industry	U.S., multi-industry	Global, financial services
% Financial services	24%	n/a	8%	11%	21%	n/a	100%
% Govt. /Public Sector	8%	n/a	9%	6%	4%	n/a	
% Other	68%	n/a	83%	83%	75%	n/a	
% revenues < \$US1 Bn	42%	0%	54%	47%	50%	n/a	51% <sup>5</sup>
ERM Development <sup>6</sup>							
% with ERM in place	11%	20%	4%	12%	30%	7%	35%
% Developing ERM	38%	14%	22%	36%	} 42%	40%	32%
% Considering ERM	42%	36%	47%	23%		19%	18%
% Not considering ERM	22%	31%	27%	29%	28%	34%	15%

Notes to Table III.2:

1. Respondents for all surveys were senior executives, e.g. CEOs, CFOs, CROs
2. The 2008 Marsh and RIMS survey report did not provide details of the survey population. It would, however, be reasonable to expect a similar population make up as for the previous Marsh and RIMS surveys.
3. 'Global' means the companies surveyed were located internationally; 'U.S.' means the companies surveyed were based in the United States. The 2007 Marsh & RIMS survey was primarily U.S. focussed but included a small percentage (6%) of firms based elsewhere.

4. The percentage of companies with multi-national operations was not consistently reported across the surveys.

5. Refers to Assets > \$US1 billion, not revenue

6. Although each of the surveys employed the term 'Enterprise Risk Management', not all of them defined it, and none of them defined what constituted 'full' vs 'partial' ERM development. Given the lack of definition, and potential for variation in respondents's perceptions of their own firms, the results are only indicative rather than definitive.

**Table III.3.** Comparative summary of industry surveys showing penetration of the role of CRO amongst global corporates

Survey	Year	% firms with CRO position	Survey scope <sup>1,2</sup>	Reference
ERM Trends & Emerging Practices	2001	24%	130 respondents, global, multi-industry, 24% financial services	(Miccolis et al. 2001)
The evolving role of the CRO	2005	45% <sup>3</sup>	137 respondents, global, multi-industry, 57% financial services	(Economist Intelligence Unit 2005)
Companies on Risk	2006	35%	441 respondents, global, multi-industry, 15% financial services	(Ernst & Young 2006b)
Companies on Risk: A/NZ Supplement	2006	70% <sup>4</sup>	54 respondents, Australia & New Zealand, multi-industry	(Ernst & Young 2006b)
Deloitte Global Risk Management Survey	2002	65%	Global, 100% financial services industry	(Deloitte and Touche LLP 2004)
Deloitte Global Risk Management Survey	2004	81%	162 respondents, global, 100% financial services industry	(Deloitte and Touche LLP 2004)
Towers Perrin Tillinghast	2006	43%	204 respondents, global, 100% insurance industry	(Towers Perrin Tillinghast 2006)

Notes to Table III.3:

1. Respondents for all surveys were senior executives, e.g. CEOs, CFOs, CAEs, CROs, corporate counsels.
2. 'Global' means that the companies surveyed were located around the world; 'Australia & New Zealand' means that the companies surveyed were based in Australia and New Zealand.
3. A further 24% of companies indicated that they planned to appoint a CRO within two years. The relatively high rate of firms with a CRO role, compared to other surveys, may be a result of the large percentage of financial services firms in the sample population of this survey.
4. Ernst & Young noted that the percentage of CROs amongst Australian and New Zealand firms was significantly higher than the international average. Although they subsequently commented that "companies in Australia and New Zealand are more progressed in their approach to risk than their global counterparts" (Ernst & Young 2006b, Australia/New Zealand Supplement, p ii), the small size of the Australia/New Zealand sample may also be partially responsible for the abnormal result.

## Professional representation of Chief Risk Officers

A multitude of institutions have emerged to stake claim to and support risk managers as a professional group, particularly in the financial services and insurance industries. At the international level, these include: the Risk Management Association (RMA), the Association of Insurance and Risk Managers (AIRMIC), the Institute of Risk Management (IRM), the Global Association of Risk Professionals (GARP), the International Association of Risk and Compliance Officials (IARCP), and the Professional Risk Managers International Association (PRMIA), to name a few. Internationally, there is also the Society for Risk Analysis (SRA), and in New Zealand, the New Zealand Society for Risk Management (NZSRM), both of which have remits much broader than just financial or enterprise risk management. The SRA, for example, has subject matter specialty groups on biological and ecological risk assessment, dose response, economic risk, risk in engineering and infrastructure, and risk policy and communication. Some of these groups are full professional associations or institutes offering university-based education, examination, and certification, while others would be better described as networks, providing their memberships with discussion forums, publications and resources, and education for continuing professional development (e.g. conferences, roundtables, and seminars).

While practising CROs may belong to one or more professional associations, depending on their backgrounds, only two organisations, of the network variety, have emerged to specifically represent the interests of Chief Risk Officers as distinct from the broader group of generic “risk professionals”. These are the Committee of Chief Risk Officers (CCRO, [www.ccro.org](http://www.ccro.org)), which serves the global natural gas and power industry, and the CRO Forum ([www.croforum.org](http://www.croforum.org)), comprised of Chief Risk Officers from major European finance and insurance companies. There are also two web portals of relevance to CROs. Compliance LLC, an international provider of risk and compliance training, provides information on the CRO role and the training and certification products offered by the company (at [www.risk-officer.com](http://www.risk-officer.com)). There is also the Enterprise Risk Management Initiative ([mgt.ncsu.edu/erm/](http://mgt.ncsu.edu/erm/)) run by North Carolina State University which offers resources, articles, research, and news on ERM, as well as links to ERM-related events (e.g. conferences, courses, etc.).

Thus, Chief Risk Officers do not yet have their own independent professional representation. Rather, representation for CROs continues to be provided by the broad range of associations representing risk management and compliance professionals more generally. It is also worth noting that other professional groups have staked claims with respect to the

task of implementing Enterprise Risk Management, including internal auditors, represented by the Institution of Internal Auditors (Mikes 2010), and management accountants, represented by the Institute of Management Accounts (Shenkir and Walker 2006, 2007). Since, as Mikes (2010, p 74) noted “CROs come from many walks of life, including internal audit, external audit, financial management, business management, and consulting”, there is considerable latitude for a variety of existing professions to lay claim to the CRO role.

## **Review of practical literature on the implementation of ERM and the CRO role**

There is a substantial trade and practitioner literature dealing with the implementation of ERM and the role of the CRO, particularly from the financial services and insurance industries. This includes:

- A growing number of characteristically similar “framework” documents, which generically describe the main objectives and process components of an ERM programme (e.g. the COSO framework, AS/NZS 4360 etc.), and a related literature which describes the characteristics of those components at different levels of maturity (i.e. ERM Capability Maturity Models; see for instance MacGillivray et al. 2007a, 2007b; OCEG 2007; RIMS 2006; RMRDPC 2002).
- A broad practice-oriented “implementation” literature which offers guidance, either for the implementation of ERM specifically, or for the organisation of risk management in organisations more generally (for example, Bowden, Lane, and Martin 2001; Frame 2003; Koller 2005; Lam 2003; Merna and Al-Thani 2005; Shenkir and Walker 2006, 2007; Ward 2005; Waring and Glendon 1998; also see Fraser et al. 2008, pp 85-88 for a list of 88 of the most important ERM-related publications globally).
- Intra- and inter-industry surveys such as those summarised earlier, along with current-issue pieces published in various professional practice journals and industry magazines, further provide an ongoing commentary on the state and trends of ERM development (for example, Beasley et al. 2009; Buehler, Freemand, and Hulme 2008; Champion et al. 2009; Fraser and Simkins 2007; Schanfield and Helming 2008), and on the evolution of the role of Chief Risk Officer (for example, Atkinson 2007; Ciccarelli 2001, 2003; Lam and Kawamoto 1997; Lam 2001; Mikes and Townsend 2007; Wood

2002).

- There is also a very large consultant-based industry offering professional services, advice, and commentary on various aspects of ERM implementation and the role of CRO (for example, Coffin 2009b; Economist Intelligence Unit 2005, 2009; KPMG LLP 2009; PriceWaterhouseCoopers 2002, 2009).

While this literature continues to grow in breadth, it remains quite shallow in terms of the level of detail and insight provided. It does, however, paint a picture of the general nature of the work in which CROs engage, and the attributes and skills necessary to be a CRO (see Tables III.5 and III.6; sources are listed in Table III.4). The nature of the CRO role is revealed through surveys, reports, and commentary on the functions, tasks, and roles that the CRO fulfils (i.e. what a CRO does or should do), the dispositions and attributes that someone aspiring to the position of CRO should have (i.e. what kind of person makes a good CRO), and the experiences of actually performing the role (i.e. how different people approach it and what they learnt along the way).

The functional aspects of the CRO role (Table III.5) point up the fact that although internal control of the organisation is a primary objective of ERM, the role of the CRO is not primarily about command and control dictums (Julien and Reiger 2003) and enforcing strict adherence to rigid procedure, but rather is about influencing and changing the behaviour of people. As Power (2005b, p 141) notes, CROs are variously described as co-ordinators, advisors, strategists, analysts, synthesists, catalysts for change, developers of best practice, designers and communicators, but *not* implementers. In other words, CROs do not own or manage risk (Roberts 2006). Chief Risk Officers commonly have responsibility for setting risk management policy, and may even have significant power in risk allocation decisions, but the responsibility for owning and managing risk ultimately lies with the Chief Executive, and, by delegated authority, with executive and line management of the organisation; and this responsibility cannot be abdicated to the CRO. The CRO is therefore not an independent manager of risk who seeks to relieve management of their responsibilities, but is, rather, responsible for ensuring the quality, comprehensiveness, and transparency of the processes and practices in which they engage in order to manage risk (Petit 2006; Power 2005b, 2007; Roberts 2006). In essence, the CRO is a 'risk governor' charged with ensuring that the organisation follows good practice and that the risk management programme is comprehensive in its coverage.

The nature of the role requires someone with strong interpersonal skills who can build relationships with actors from the Board down to the worker's on the shop floor, and across

the breadth of the organisation's functions (Table III.6). Since the CRO is primarily an integrator and co-ordinator, the role requires someone with a breadth of knowledge and experience across risk domains, and who is comfortable working in a multi-disciplinary environment. The CRO must be able to bring together people and information from a variety of sources, facilitate collaborative knowledge production, and then synthesise and communicate that knowledge to others. The ability to integrate information from disparate sources, and to step back and 'see' the strategic implications at the enterprise level (i.e. the big picture) is regarded as essential.

**Table III.4.** Sources for Tables III.5 and III.6

Table III.5	(Aabo et al. 2005; AIRMIC, ALARM, and IRM 2002; Brown 1997; Conference Board of Canada 2001; COSO 2004; de la Rosa 2007; Gordon-Clark 2007; Graziano and Aggarwal 2005; Julien and Rieger 2003; Lam 2000, 2001, 2003; Lam and Kawamoto 1997; Lamser and Helland 2000; Lee 2000; Lee and Shimpi 2005; Marsh and RIMS 2004, 2005; Mikes and Townsend 2007; O'Rourke 2007; Roberts 2006; Salvador 2007; Shenkir and Walker 2006; Utter 2006; Walker and Shenkir 2008; Ward 2001)
Table III.6	(Atkinson 2007; Ciccarelli 2001, 2002, 2003; Conference Board of Canada 2001; de la Rosa 2007; D'Elia 2002; Economist Intelligence Unit 2005; Eick 2003; Falkenstein 2001; Gordon-Clark 2007; Kelly 2007; Lam 2001; Lee 2000; Lee and Shimpi 2005; Louisot 2003; Marsh and RIMS 2004; Woodrow 2005)

**Table III.5.** Common functional aspects of the role of CRO  
(sources for this table are listed in Table III.4, above)

<p><i>The nature of the CRO's work varies with time and place, and while any given performance of the role in a specific time and place in the organisation may emphasise certain functions over others, any such performance necessarily involves all of the aspects in this table to some degree. In this regard, the functions listed here are not discrete activities in which the CRO engages, but rather should be understood as constitutive aspects of a complex performance.</i></p>	
Functional Aspect	Description
Framework development	Developing and establishing an integrated risk management framework. Developing risk policies and defining risk appetite/tolerances. Establishing and enforcing risk management governance processes, roles and accountabilities.
Leadership	Ultimate champion of risk management framework and processes, providing overall leadership, vision, and direction for risk management. Cultivating Board and senior management support for risk management. Setting strategy for risk management development. Chair (or member) of Risk Management committee.
Cultural development	Developing a risk-aware culture, and deeper understanding of risks and their inter-relationships throughout the organisation. Developing and implementing risk management training programme. Establishing risk-based performance measurement and incentives. Extending risk management principles into wider business processes and strategy.
Communication	Developing and communicating an integrated picture of risk across the enterprise. Improving the representation and communication of risk within the organisation, and to the Board and stakeholders. Reporting to Board and stakeholders.
Co-ordination	Centralising and co-ordinating an organisation's risk management efforts. Monitoring the enterprise risk portfolio. Approval of key risk management decisions. Formal oversight of risk management efforts. Assessment and monitoring of the application and effectiveness of risk management processes.
Infrastructure development	Implementing risk metrics and reports. Developing analytical systems and data management capabilities to support risk management. Developing and disseminating risk identification and assessment methodologies and tools. Preparation and dissemination of manuals, guidelines, and procedures.
Facilitation	Facilitates or otherwise assists with risk identification and assessment processes within the organisation. Acts as an advisor to the business. Acting as a conduit for the interchange of information on risk and risk management.
Insurance management	Oversight of insurance programme. Identification of insurance needs and negotiation with insurance brokers.



**Table III.6.** Attributes and skill sets of effective CROs  
(sources for this table are listed in Table III.4, above)

Attribute	Description
Highly educated with broad spectrum of knowledges and experiences	<p>CROs are highly educated, holding at least a bachelors degree along with some form of continuing professional education accreditation. More advanced degrees are also common (some surveys report up to 50% of risk managers holding Masters or PhD degrees). The vast majority also have more than five years experience in the risk management field. While some technical and quantitative skills are necessary, the CRO is more appropriately a generalist with knowledge and experience across a broad spectrum of areas:</p> <ul style="list-style-type: none"> <li>▪ Financial (credit and market risk, and risk financing methods, and knowledge of financial and accounting processes)</li> <li>▪ Insurance (insurance instruments, and risk transfer mechanisms)</li> <li>▪ Legal (environmental, statutory, contract, case and tort law, as well as knowledge of judicial process)</li> <li>▪ Risk management process</li> <li>▪ Industry specific (especially with regard to sources of operational risk)</li> <li>▪ General business management</li> </ul>
Integrative, big-picture thinker	<p>Ability to quickly develop detailed understanding of business operations and the relationships between different parts and levels of the business.</p> <p>Ability to parse and integrate large amounts of information from different sources, assess patterns, and grasp the 'big picture' (i.e. strategic implications).</p> <p>Comfortable working in a multi-disciplinary environment.</p>
Strong communication skills	<p>Ability to understand the decision-making needs of different groups throughout the organisation, and to organise and communicate relevant risk information to those groups (including the Board and external stakeholders, senior and line management, and workers 'at the coal-face').</p>
Strong leadership, co-ordination, and facilitation skills	<p>Ability to lead disparate groups through the risk management process.</p> <p>Ability to bring together coalitions and facilitate collaborative work on common risk issues.</p> <p>Ability to facilitate knowledge production processes with functional-specific, and cross-functional groups in the organisation.</p> <p>Ability to co-ordinate risk management activities across the organisation.</p>
Strong political acumen and interpersonal skills	<p>Ability to understand and negotiate the political landscape.</p> <p>Ability to interact, and build relationships and partnerships with people at different levels throughout the organisation.</p> <p>Ability to influence people to get things done.</p>

# Review of academic literature on the implementation of ERM and the CRO role

## On Enterprise Risk Management

### **Harrington et al. (2002) The United Grain Growers Case**

Strictly speaking, the paper by Harrington et al. (2002) was not about the implementation of ERM as ERM is currently conceived in world-level standards. Rather the paper described the process (essential brainstorming and analysis guided by an external consultant) that the company went through to identify the major source of risk (weather volatility affecting grain yield) contributing to its earnings volatility, and how the company addressed this risk through an innovative insurance contract. In this regard, the United Grain Growers case should really be considered an example of strategic risk assessment, rather than a study of the implementation of integrated, enterprise-wide risk management within the firm.

### **Aabo et al. (2005) The Hydro One Case**

Aabo, Fraser, and Simkins (2005) described the implementation of ERM at Hydro One, a large electricity delivery company in Ontario, Canada. The paper covered the creation of the CRO position and the Corporate Risk Management Group in 1999, the development of the ERM policy and framework, a pilot risk identification workshop carried out to identify the ten most critical risks facing the company, the processes and tools that the company employs to identify and assess risks, an overview of the corporate risk profile and how risk information feeds into the asset management (resource allocation) process, and the benefits of ERM to Hydro One (Aabo et al. 2005). However, while the paper provided some useful insights, particularly about the use of various tools (e.g. workshops, the Delphi Method, risk maps) to create the enterprise risk profile, it was necessarily a very high level overview of a complex process that evolved over five years, and thus did not provide the kind of detailed “how to’s” called for by Fraser et al. (2008). Although co-authored by two academics, the paper did not move beyond description to the generation of theoretical insight.

### **Stroh (2005) The UnitedHealth Group Case**

Stroh’s (2005) paper covered similar ground to the Hydro One case, but with more emphasis on the objectives, drivers, and success factors of the implementation process. Again, however,

the paper was necessarily a very brief summary of a long and complex organisational change process. The paper identified important themes which are commonly emphasised in practical literature on ERM (e.g. creating risk awareness, alignment with strategy, securing strong executive backing etc.), but did not provide any specific guidance on how these things were achieved within the company. Contrary to its inclusion in Iyer et al.'s (2010) review of *academic* research on ERM, the paper does not present academic research nor tie the practical insights provided back to relevant financial or management theory. The publication (*Strategic Finance*) in which the case appeared is a professional practice magazine published by the Institute of Management Accountants, not a peer-reviewed academic journal, and the author of the case did not write in an academic capacity (at the time of publication, he was the director of Business Risk Management at UnitedHealth Group).

### **Nocco and Stulz (2006) Theory and Practice of ERM**

Nocco and Stulz's (2006) paper presented a corporate finance theory perspective of ERM, but was not strictly a case study. Rather, in collaborative effort between academic (Stulz) and practitioner (Nocco, CRO of Nationwide Insurance), the paper combined corporate finance theory with illustrative references to risk management practices at Nationwide Insurance to present a theoretically grounded conceptual framework for ERM implementation. The paper covered how ERM creates shareholder value, at both a "macro" or company-wide level and a "micro" or business-unit level, and how a company can determine the optimal amount of total risk to bear, before discussing various aspects of ERM implementation in a firm. Of the five case studies reviewed by Iyer et al. (2010), this was the only one to succeed in bringing relevant theory to bear on the task of describing ERM. However, the paper reflects a strong corporate finance perspective, grounding the ERM approach in concepts such as earnings volatility, credit ratings, financial distress, Economic Capital, and Value-at-Risk.

### **Acharyya (2008) ERM in European Insurers**

Acharyya (2008) reported a study of ERM practices at four Europe-based multi-national insurance companies. The study sought to investigate the understanding, evolution, design, and performance of ERM in the different companies, and the challenges they faced while implementing ERM. Data was collected from structured interviews with participants in two companies, and from a survey of participants in the other two companies. The key findings were: (1) company personnel variously interpreted ERM as a process, a tool, or an approach, depending on how they encountered the risk management function in their roles; (2) ERM

implementation was driven primarily by CEO perceptions and attitudes towards various contextual factors, of which the most important was regulatory compliance; (3) the four companies reported a variety of operational and technical challenges to implementing ERM, prominent among which were the development of a common language and culture of risk, the measurement of operational risk, and the accuracy, consistency, and adequacy of data; and (4) that the implementation of ERM was complex, time consuming, and costly, that ERM took on different forms in different contexts, that there was no rigorous framework for evaluating the performance of ERM, and that managers reported a variety of potential benefits of ERM, none of which could be concretely proven. The paper also described a generic five stage conceptual model of ERM implementation drawing on insights from the study, but, like the above framework from Nocco and Stulz (2006), the model is grounded in the calculation of financial constructs (i.e. Economic Capital), which reflects the insurance industry context of the study.

- **Note:** In Iyer et al. (2010), and elsewhere, the reference to this study is Acharyya and Johnson (2006), reportedly published in a special issue of *The Geneva Papers on Risk and Insurance: Issues and Practices*, July 2006. However, no such paper is listed in the PalgraveMacmillian archive for this journal(\*), or in the contents of special issues listed on the Geneva Association website(\*\*), including the July 2006 issue in which the paper was claimed to have been published. I have therefore chosen to cite the later paper by Acharyya (2008).

(\*<http://www.palgrave-journals.com/gpp/archive/index.html>)

(\*\*[http://www.genevaassociation.org/Publications/Geneva\\_Papers\\_on\\_Risk\\_and\\_Insurance.aspx](http://www.genevaassociation.org/Publications/Geneva_Papers_on_Risk_and_Insurance.aspx))

### **Mikes (2005) ERM in Action**

Mikes (2005) drew on two empirical cases of ERM implementation to characterise ERM in practice as reflecting a mix of risk management practices particular to an organisation and its context. The paper suggested that the “risk management mix” in any organisation can be described as a combination of four “ideal types”: Risk Silo Management, Integrated Risk Management, Risk and Value Management, and Strategic Risk Management. The characteristics of the four “ideal types” were drawn from four themes emphasised in the ERM literature (respectively, risk quantification, risk aggregation, risk-based performance measurement, and the management of non-quantifiable risks), with each type differentiated by institutional background, calculative focus, and associated calculative techniques and

tools. The empirical work described the particular mix of ideal types observed at two banks, and suggested that in each case the mix was a function of certain contextual (firm-specific and institutional) factors. Mikes paper stands out because of the methodological approach (70 in-depth interviews with participants at the two banks) and the level of detail provided by the narrative style, including liberal use of direct quotes from participants. Because Mikes targeted her research at “the coal-face” of ERM practice she was able to show how the performance of ERM in action is intricately tied up with the personalities, attitudes, and culture of the organisation, the pragmatics of doing business, and the idiosyncracies of corporate history. Rather than attempting to explain ERM in terms of financial theory, she drew upon concepts from organisation and management, and sociology to make sense of what was going on in the two banks.

## On the role of Chief Risk Officer

Academic research on the role of Chief Risk Officer is practically non-existent. To my knowledge only three authors have published academic research on the role. Ward (2001) published an early study describing a range of tasks, objectives, and drivers for the CRO role. Power (2005b, 2007) provided a discussion of the CRO role as a recent addition to an existing category of “regulatory officers”, and as a new professional role. More recently, based on a study of 15 CROs, Mikes (2007; 2010) described three archetypal CRO roles.

## The activities of Corporate Risk Managers (Ward 2001)

Ward (2001) published what may be the earliest academic research on the CRO role, an exploratory study of the activities of corporate risk managers (the British term for CROs) in thirty organisations. The study revealed six major functions in which CROs were typically engaged, nine objectives to which CROs aspired, and five major contextual factors affecting the direction and extent of the CRO role (Table III.7).

*Table III.7. Functions and objectives of the CRO role, and factors influencing the direction and scope of the role (from Ward 2001)*

Functions	Objectives	Influencing factors
<ul style="list-style-type: none"> <li>▪ Designing and establishing an integrated risk management strategy and policy.</li> <li>▪ Establishing and maintaining a detailed RM methodology.</li> </ul>	<ol style="list-style-type: none"> <li>1. Limit potential financial losses and liabilities.</li> <li>2. Minimise the total cost of risk.</li> <li>3. Manage risk at the lowest</li> </ol>	<ul style="list-style-type: none"> <li>▪ Top management influence.</li> <li>▪ External influences (esp. regulatory requirements).</li> </ul>

*Table III.7. Functions and objectives of the CRO role, and factors influencing the direction and scope of the role (from Ward 2001)*

Functions	Objectives	Influencing factors
<ul style="list-style-type: none"> <li>▪ Insurance management.</li> <li>▪ Providing advice to business units on RM techniques.</li> <li>▪ Monitoring the application and effectiveness of RM processes.</li> <li>▪ Co-ordination of the delivery of information on risk and RM.</li> <li>▪ Acting as a conduit for the interchange of information on risk and risk management.</li> </ul>	<p>possible cost.</p> <ol style="list-style-type: none"> <li>4. Maximise return on investment for risk management effort.</li> <li>5. Manage and reduce the cost of insurance.</li> <li>6. Try to remove the need for insurance.</li> <li>7. Protect and preserve both tangible and intangible assets from loss or damage.</li> <li>8. Make risk management part of the decision process.</li> <li>9. Establish, analyse, and control the organisations risk profile.</li> </ol>	<ul style="list-style-type: none"> <li>▪ Nature of the business (size, industry, culture).</li> <li>▪ Corporate developments (growth, history).</li> <li>▪ Characteristics of the risk management department (experience, resources).</li> </ul>

### **The CRO as a new professional role (Power 2005b)**

Power (2005b) explicitly links the role of CRO to the implementation of Enterprise Risk Management in organisations, and interprets the role by translating Abbott's (1988) notion of competing professions to the organisational level. For Power (2005b, 2007), ERM frameworks are rational designs for risk management systems which reconceptualise the organisation from a risk management perspective, and the Chief Risk Officer is that corporate role explicitly concerned with operationalising these designs within organisations.

As the ultimate champion of ERM, the CRO is charged with internalising world-level ideas about what constitutes 'good practice' for the governance and control of organisations (Power 2005a, 2005b). In this role the CRO seeks to rearrange and reposition existing tasks and routines under the new umbrella of ERM, effectively competing "to have a particular organizational model accepted internally" (Power 2005b, pp 143-144). Power describes the role as a recent addition to the broader category of "functionally dedicated 'officers' with some degree of regulatory role inside organisations", positions which "provide internal organizational representations of externally encountered norms and rules" (2005b, p 137). The creation of an officership role signals "organizational seriousness about issues", and actors in this group (which includes compliance officers, health and safety officers, environmental risk officers, race relations officers) are tasked with articulating the business case for compliance, promoting changes in management and control practices, and

mobilising other internal actors with 'win-win' rhetoric (Power 2005b, pp 137-139). The history of such roles, however, suggests that they are difficult to fulfil, encountering complex organisational politics, constant challenges to their effectiveness and legitimacy, and jurisdictional conflicts with other management roles (Power 2005b, 2007). It is a position where CROs must work constantly to maintain the legitimacy, functionality, prestige, and value of ERM (Power 2005b).

### **Archetypal roles of the CRO (Mikes 2010)**

While the specific focus and scope of the risk management function and of the CRO's mandate may vary widely from context to context (Ward 2001), two distinct types of CRO have become visible, identifiable by the primary focus of the role (Atkinson 2007).

The first, reflecting the strong regulatory heritage of ERM, is that of the CRO as "chief compliance officer", responsible for regulatory compliance and reporting (e.g. with Sarbanes-Oxley and Basel II). This form of the CRO role, common in highly regulated industries such as financial services and pharmaceuticals, is focussed primarily on ensuring that the organisation has the necessary formal processes in place for compliance with the increasing number of regulatory obligations imposed by governments and other external agencies (Lee and Shimpi 2005). Indeed, regulatory compliance is consistently ranked by executives as a high priority for their risk management programmes (Deloitte and Touche LLP 2007; Economist Intelligence Unit 2005; Ernst & Young 2006a, 2006b; KPMG International 2007).

The second type of role is that of the CRO as "a genuine partner in business management" playing an important role in strategy setting and decision making (Atkinson 2007, p 26). This latter version of the role, where the CRO acts as business advisor reflects the value creation proposition of ERM and the change in emphasis in risk management from internal control to decision enhancement (Lee and Shimpi 2005; Wood 2002).

More recently, drawing on a study of 15 CROs, Mikes (2010) extended the above categorisation, describing three archetypal CRO roles (CRO as compliance champion, CRO as strategic controller, and CRO as strategic advisor), differentiated by the degree of emphasis placed on compliance with regulatory and risk management standards, the extent and sophistication of a firm's risk modelling, and the attitudes of individual CROs. The first of these roles reflects the traditional focus on regulatory compliance, and entails "advocating and policing compliance with pressing stakeholder requirements and keeping up with new regulations and standards" (Mikes 2010, p 75). Of more interest with respect to this thesis, however, is Mikes' description of the latter two roles:

- **CRO as strategic controller:** building on the availability of firm-wide risk models, this role approaches the risk function as “a formal risk-adjusted performance management system” (Mikes 2010, p 77). Chief Risk Officers in this category “preside over the close integration of risk and performance measurement and ensure that risk-adjusted metrics are deemed reliable and are relied on. They advise top management on the absolute and relative risk-return performance of various businesses and influence how capital and investments are committed” (Mikes 2010, p 77).
- **CRO as strategic advisor:** in this role, CROs “command board-level visibility and influence, predominantly as a result of their grasp of emerging risks and nonquantifiable strategic and operational uncertainties. They bring judgement into high-level risk decisions, challenge the assumptions underlying business plans, and use traditional risk controls and lending constraints to alter the risk profiles of particular businesses” (Mikes 2010, p 78).

The two roles are differentiated by the purpose of and degree of reliance on quantitative risk modelling, reflected in the attitudes of different CRO-types toward quantification. Chief Risk Officers of the “strategic controller” variety exhibited what Mikes referred to as *quantitative enthusiasm*, characterised by a commitment to extensive risk modelling, a belief in such models as robust and relevant tools in decision making, and a primary objective to measure the aggregate risk profile of products and business lines (Mikes 2010). In contrast, CROs of the “strategic advisor” variety exhibited what Mikes referred to as *quantitative skepticism*, characterised by a belief that risk modelling is “not sufficiently accurate to produce an objective picture of the underlying risk profiles” and that such quantitative calculations should be used as trend indicators only (Mikes 2010, p 78). The emphasis in the “strategic advisor” role is on playing “devil’s advocate” and facilitating the cross functional sharing of risk information to prevent “risk incubation” and to enhance “risk anticipation” and learning about risk profiles, particularly for non-quantifiable uncertainties (Mikes 2010).

Mikes suggests that the distinction (which is not black and white) between the two roles ultimately rests on a philosophical choice over “where to draw the line between what can be reliably measured and modeled and what must be placed in the hands of qualitative judgement” (2010, p 79). Both roles also require very different capabilities. The “strategic controller” role “calls for building a sophisticated risk-modelling capability which is foundational to risk-based performance measurement”, while the “strategic advisor” role “requires an intimate knowledge of the business and what can go wrong – experience that risk officers can only gain by having lived through many organizational successes, losses, and



crises” (Mikes 2010, p 80).

Notably, Mikes also calls attention to the fact that, despite the aspirations of “strategic controllers” to quantitative objectivity, the accuracy of predictions about what might happen is always undeniably limited, while the construction of risk-adjusted performance measures is inherently political, being both the product of processes of organisational consensus and affecting resource and reward allocations. As such CROs must be “modest in their claims of objectivity” (Mike 2010, p 78).



# Appendix IV

## Proposals to improve the quality of Watercare's "risk data"

### **Concerns about Watercare's "risk data"**

The CRM quickly developed an overriding concern with what he perceived to be the poor quality of the company's "risk data". As he used it, the term "risk data" referred to the informational outputs of risk assessments performed by employees, which were formally recorded in the company's corporate risk register as statements of risk in a standardised format (i.e. cause, consequence, likelihood, severity, controls). In the CRM's opinion the data in that register exhibited a number of problematic symptoms: (i) a "high degree" of subjectivity and ambiguity in risk descriptions; (ii) incomplete, overlapping, or duplicated risk descriptions; (iii) a lack of detail and transparency in risk assessments; (iv) meaningless enterprise risk profiles; and (v) a lack of structure.

### **Risk descriptions**

The CRM and the external auditors identified that the risk descriptions and assessments in the corporate risk register were often ambiguous, incomplete, overlapping, or simply too brief (Clement 2007e; PriceWaterhouseCoopers 2007). Entries in the register often described only the cause or the consequence of risks without clearly linking the two together, or lumped together multiple causal factors, each with different likelihoods, potentially different consequences, and different controls, in a single risk description. For example, the failure of a

dam embankment could result in a catastrophic breach of the dam, release of the impounded water, extensive downstream damage, and significant financial and operational impacts on the company, but in the risk register, a dam embankment failure was identified as a single risk, with low probability and high consequence (extract from Watercare Risk Register, February 2005). That single risk entry specified multiple potential causal factors (internal erosion, overtopping during probable maximum flood, structural failure in a seismic event, and blockage of the bellmouth spillway), each of which was independent (the likelihood of the probable maximum flood is orders of magnitude higher than the likelihood of a major seismic event, and both are potentially less likely than a spillway blockage), and each of which had different management controls (the possibility of internal erosion depends on the condition of the embankment underdrains, while the emergency spillway would discharge excess water in the event of a blockage of the main spillway).

## **Transparency**

The ambiguity and subjectivity of risk assessments was further compounded by the fact that the risk register was not transparent as to the quality of the assessment behind each risk entry (Clement 2007e). Risk identification was often unstructured, the most common methods being brainstorming workshops or informal suggestions, while the assignment of consequence and likelihood scores usually gauged by qualitative judgement (i.e. selecting qualitative consequence descriptions from the risk framework on the basis of subjective or intuitive expectations). The problem was that the entries in the risk register did not record whether the final risk score was the result of an exhaustive, quantitative analysis, or simply a quick 'back of an envelope' assessment of someone's gut feeling.

This was quite clearly revealed in the treatment of inherent and residual risk scores in the register, which were intended to reflect the effect of risk controls ('inherent' risk being, conceptually, the risk that would exist without any controls, and 'residual' the risk that would exist with all controls functioning as intended)<sup>27</sup>. Although various risk controls were usually listed for each risk in the register, there was no transparency in how those controls translated into adjustments made to the residual risk score (Clement 2007e). Indeed, the residual risk score sometimes reflected a reduction in consequences (or likelihood) even where the

---

<sup>27</sup> In principle, the difference between the inherent and residual risk scores should provide a measure of the value of the designated risk controls. However, while the concepts of inherent and residual risk are specified in AS/NZS 4360, they are difficult to operationalise in practice because they represent opposite ends of a scale of 'pessimism or optimism' about the reliability of the implemented risk controls.

specified controls would have no impact on the consequences (or likelihood) of the risk. The lack of transparency in such qualitative assessments was exacerbated by the subsequent reduction of the five consequence ratings to a single root-mean-squared (RMS) average risk score, which, although it simplified risk reporting (one risk, one number), also further obscured the underlying information (Clement 2007e).

### **Meaningless Enterprise Risk Profiles**

The CRM was also concerned about the mathematically incorrect representation of cumulative risk profiles at the enterprise level, which he saw as a consequence of the qualitative and subjective nature of the underlying risk analyses. The format for monthly risk reporting to the Chief Executive and the Board was to present the total number of risks in each business group and their distribution between the five risk classes. A similar representation, broken down by asset type, can be found in Watercare's asset management plans prior to 2006 (e.g. Watercare Services Ltd. 2005a, p 36). Similarly, the Enterprise Risks reported to the Board were really just groupings of risks associated with various critical assets, the Enterprise Risk score being simply the highest individual risk score from that group (Clement 2007e).

This form of representation was problematic for a number of reasons: the number of risks in each category is dependent upon the comprehensiveness and depth of risk identification and analysis undertaken; each risk is counted equally in the final summation which effaces the considerable variation in the depth and quality of each risk assessment; and it tells decision makers nothing about the cumulative organisational exposure to different types of consequences across all risk classes and categories, nor about the significance of having a great many Class 3 risks relative to a single Class 5 risk (especially where the Class 5 risk may be high consequence, low probability). The result was that the enterprise risk profiles were effectively meaningless: "...those numbers mean nothing, it's like throwing a handful of darts at a dart board and saying there's your profile" (CRM commenting on the crudeness of the enterprise risk profiles, 15 June 2007).

### **Lack of structure in the risk register**

While there was no formal restriction on the scope of the corporate risk register, it contained an overwhelming number of asset-focussed risks, typically each risk being defined as the loss or failure of a specific asset or group of assets (Clement 2007e). Although one could argue that this operational bias simply reflected the fact that operating physical assets was, in the main,

what Watercare did, the CRM felt that the content of the register was more a consequence of way the register had evolved historically, driven from the bottom-up by Operations staff who populated the register with their ‘worries’ (notes from interview with CRM, May 2007).

The CRM argued that the lack of strategic focus and hierarchy in the register was problematic for two reasons. It meant that the enterprise risk profile was skewed and did not adequately reflect other important activities (and risks) across the company, and, since the entire body of corporate risks (~800) was contained within a single undifferentiated layer, there was no way to assess interdependencies between risks, or to generate meaningful cumulative risk profiles from the existing data set (Clement 2007e). These issues were also identified in the PWC audit report (PriceWaterhouseCoopers 2007).

## **Proposals to improve data quality**

The following pages (Table IV.1) present the Corporate Risk Manager’s various proposals for improving the quality of Watercare’s “risk data”, organised according to common functions. The proposals are drawn from his Risk Management Framework Development Plan of March 2007 (Clement 2007e), and his Risk Management Vision of May 2007 (Clement 2007f).

**Table IV.1.** The CRM's proposals to improve data quality (collated by function)

Function	Target Functionality + Capability Development Tasks
Defining the risk management context (2007e: D1, D5, D19) (2007f: s3.1)	<p><u>Hierarchy of objectives:</u> The risk management framework should support a set of explicit WSL objectives. Sub-dividing and cascading the organisations strategic objectives through the various levels of management (as per the Balanced Scorecard Concept) will provide a basis for identifying risks appropriate to each management level.</p> <ul style="list-style-type: none"> <li>▪ <i>Develop suitable risk scoring frameworks for each business context.</i></li> <li>▪ <i>As initial priority, develop appropriate risk scoring framework for the project management context.</i></li> </ul>
Improving risk identification capabilities (2007e: D2, D3, D9, D18) (2007f: s2.3/3.4, s2.5/3.6)	<p><u>Adoption of tools and methods:</u> The employment of recognised risk identification tools and methods would enhance the probability that a comprehensive set of risks are identified.</p> <ul style="list-style-type: none"> <li>▪ <i>Establish risk identification methodologies for different business contexts.</i></li> <li>▪ <i>Conduct risk identification workshops.</i></li> </ul> <p><u>Clarity of risk descriptions:</u> The risk descriptions (in the risk register) need to be unmistakable. Accordingly risks should be defined as impact against a specific WSL objective due to a precise cause.</p> <ul style="list-style-type: none"> <li>▪ <i>Establish risk description criteria, and re-describe existing risks according to these criteria.</i></li> </ul> <p><u>Knowledge capture and application:</u> Continual knowledge capture and the subsequent refinement of risk information are important in maturing the risk management framework. A formal process should be established for collating, analysing and feeding knowledge from incidents and maintenance (Asset Specialist Group) back into the risk identification and evaluation stages. In the immediate future this could be achieved by regular periodic reporting.</p> <ul style="list-style-type: none"> <li>▪ <i>Investigate and develop databases for maintaining records of incidents.</i></li> </ul>

**References** (sections cited above):

WSL Risk Management Framework Development Plan, March 2007 (Clement 2007e);

WSL Risk Management Vision, May 2007 (Clement 2007f).

**Table IV.1.** *The CRM's proposals to improve data quality (collated by function)*

Function	Target Functionality + Capability Development Tasks
Improving risk evaluation and decision making (2007f: s2.1/3.2)	<p><u>Risk tolerance criteria:</u> transparency around acceptance criteria for water and wastewater service supply risk will be particularly important in the future with the AMP budget expected to increase substantially. Risk tolerance thresholds should be explicitly defined by an array of likelihood and consequence values for each of the different contexts that is used to measure consequences in the risk management framework (i.e., financial, legal, service continuity etc).</p> <ul style="list-style-type: none"> <li>▪ <i>Review contractual and legal requirements, and national and international standards. Analyse the organisation's internal resilience in different business contexts. Develop risk tolerance criteria, and present to and consult with the Board for approval. Possible consultation with stakeholders (shareholders, ARC, WSAA etc.).</i></li> </ul> <p><u>Investment uncertainty evaluation:</u> An Investment Uncertainty Evaluation will be presented within those CAPEX applications that require Board approval. The evaluation will convey the levels of uncertainty associated with the various expenditure options by presenting the spread of possible risk reduction and the capital allocation calculations (i.e., the spread of possible cost benefit levels). It is anticipated that this will be calculated using Monte Carlo analysis... a risk aggregation methodology will be required, whereby component (source) risks can be numerically combined to calculate the overarching risk. It is anticipated that fault tree analysis techniques will be used to achieve this goal.</p> <ul style="list-style-type: none"> <li>▪ <i>Develop a methodology to aggregate risks to produce a distribution of potential outcomes vs likelihoods for a given set of risks.</i></li> <li>▪ <i>Develop an agreed capital allocation methodology.</i></li> </ul>

**References** (sections cited above):

WSL Risk Management Framework Development Plan, March 2007 (Clement 2007e);  
 WSL Risk Management Vision, May 2007 (Clement 2007f).



**Table IV.1.** The CRM's proposals to improve data quality (collated by function)

Function	Target Functionality + Capability Development Tasks
<p>Establishing enterprise-wide risk assessment capabilities</p> <p>(2007e: D4, D7, D10, D11, D13)</p> <p>(2007f: s2.2/3.3)</p>	<p><u>Risk register capabilities – aggregate analysis:</u></p> <p>The risk register requires the ability to record risks within a hierarchical structure, such that the body of risks in the register may be aggregated to form higher-level risks. This requires that the score for the higher-level risks be calculated from the likelihoods and consequences of the sub-set risks. An appropriate method of summation will be required.</p> <p>Such a structure would also allow risk causation to be traced top-down from the enterprise risks to lower-level risks throughout the organisation. It would also provide for the ability to analyse common sources of risk (e.g. potential failure causes that may be common across multiple higher level risks – potential examples include process control software failure, inadequate competency of maintenance staff, SCADA communications failure, failure of the server supporting Mosaic etc.) and common critical controls upon which the company places significant reliance (Incident Management Plans being one example).</p> <ul style="list-style-type: none"> <li>▪ <i>Develop the risk register structure.</i></li> <li>▪ <i>Develop a method to aggregate 'sub risks' into a single higher level measure, such that total cumulative risk may be calculated for any given risk profile. Establish a baseline risk profile.</i></li> </ul> <p><u>AMP formulation tool:</u> The AMP Formulation tool will support executive approval of the proposed AMP by evaluating and optimising the cumulative risk reduction on the Watercare risk profile. It is anticipated this will be achieved via quantitative risk analysis of various combinations of projects, subject to elective, financial and programming constraints. This will require modelling the future likelihood and consequences of asset failure for two possible scenarios; (1) deferring project works, and (2) implementing a proposed project. Risk profiles may need to be approximated at multiple points in the future (e.g. 5, 10, 20 years).</p> <ul style="list-style-type: none"> <li>▪ <i>Develop a methodology for estimating future risk profiles. Investigate the use of AvSim software to model future network configurations under predicted demand conditions, and to gauge asset criticality under those conditions.</i></li> <li>▪ <i>Develop a methodology to evaluate various combinations of risk control projects to determine the maximum risk reduction for any given AMP investment.</i></li> </ul>
<p>Developing a 'risk-aware' culture</p> <p>(2007e: D20, D23)</p> <p>(2007f: s2.8)</p>	<p><u>Staff training programme:</u> There is a need for ongoing risk management training programmes to maintain and develop knowledge and competency. Internal training programmes should be developed and delivered internally to cover an introduction to risk management concepts, and the WSL risk management framework.</p> <ul style="list-style-type: none"> <li>▪ <i>Develop staff training programmes.</i></li> <li>▪ <i>As a priority, develop project risk management training.</i></li> </ul>
<p><b>References</b> (sections cited above):</p> <p>WSL Risk Management Framework Development Plan, March 2007 (Clement 2007e);</p> <p>WSL Risk Management Vision, May 2007 (Clement 2007f).</p>	



# Appendix V

## Normative models of Risk Management

This appendix presents a review of key ERM concepts, expressed in the international literature (standards, frameworks, text books etc.), as a standardised “picture” of what ERM should look like. It also presents a review of key concepts in engineering risk assessment as a standardised “picture” of the engineering approach to the assessment of risk.

### **The principles of ERM**

The label of “Enterprise Risk Management” does not refer to a unified set of practices or a single standard, but rather to a category of ideas, frameworks, and process-concepts for thinking about the organisation and oversight of risk management activity on an enterprise-wide basis (Power 2005a, 2007). A number of standards, frameworks, and guidelines for ERM, have been published by various professional organisations around the world. Prominent among others are the COSO *Enterprise Risk Management – Integrated Framework* (2004), the Australia/New Zealand Standard *4360 Risk Management* (2004a), the recent International Standard *ISO 31000 Risk Management – Principles and guidelines on implementation* (2009), and *A Risk Management Standard* (2002), a joint publication by the Institute of Risk Management (IRM), the Association of Insurance and Risk Managers (AIRMIC), and The National Forum for Risk Management in the Public Sector (ALARM). In addition, there is a significant international literature, which seeks to provide guidance for the implementation of ERM in organisations (for example, Barton et al. 2002; DeLoach 2000; Frame 2003; Lam

2003; Merna and Al-Thani 2005; Meulbroek 2002; Miccolis and Shah 2000; Nocco and Stulz 2006; Shenkir and Walker 2006, 2007; Ward 2005).

Documents such as the COSO and ISO 31000 standards, and other similar frameworks, are not blueprints for action (Power 2004, 2007), but rather sketch in outline the main procedural and structural components of an ERM programme. The central component of ERM is the generic risk management process, which is typically defined in multiple phases: e.g. establish the context, identify the risk, analyse the risk, evaluate the risk, treat the risk, monitor and review (AS/NZS 4360:2004). This generic process has been codified in a common form in various international standards (Ward 2005, p 152).

## The distinguishing features of ERM

Among the world-level standards and frameworks, and within the broader literature, definitions and approaches to ERM vary, but common elements are: an approach to the identification, analysis, evaluation, and management of risk which is structured, rigorous, and systematic, with an holistic, enterprise-wide focus, and seamlessly integrated into all aspects of organisational decision making such that risk management becomes part of the organisation culture (see Table V.1).

*Table V.1. Characteristics of ERM, as defined in world-level frameworks*

<b>COSO Framework description of ERM</b> (COSO 2004, p 17)	<b>ISO/DIS 31000 description of ERM</b> (ISO/DIS 31000:2008)
<p>Enterprise Risk Management is:</p> <ul style="list-style-type: none"> <li>▪ a process, ongoing and flowing through an entity;</li> <li>▪ effected by people at every level of an organisation;</li> <li>▪ applied across the enterprise, at every level and unit, and includes taking an entity-level portfolio view of risk;</li> <li>▪ designed to identify potential events that, if they occur, will affect the entity and to manage risk within its risk appetite;</li> <li>▪ able to provide reasonable assurance to an entity's management and board of directors; and</li> <li>▪ geared to the achievement of objectives in one or more separate but overlapping categories.</li> </ul>	<p>Effective Risk Management is:</p> <ul style="list-style-type: none"> <li>▪ customised to the organisation through alignment with the entity's external and internal context (objectives, processes, structures, capabilities, and culture);</li> <li>▪ is embedded within the organisation's practices and business processes so that it is an integral part of and not separate from those processes;</li> <li>▪ is integral to decision making such that all decision making within the organisation, whatever the level of importance and significance, involves the explicit consideration of risks;</li> <li>▪ is systematic, structured, and timely, and based on the best available information; and</li> <li>▪ is comprehensive, addressing all those events that might enhance, prevent, degrade, or delay the achievement of the organisation's objectives</li> </ul>

What distinguishes ERM from isolated performances of risk management activities in specific times and places within an organisation is, first, the notion of integration (Lam 2003, p 45); in terms of (i) taking a cross-functional, enterprise-wide perspective of risk, and (ii) embedding risk management throughout the organisation and its processes; and second, the explicit centering of the organisation's strategic objectives as the focal point for all encounters with risk. The multi-dimensional nature of ERM is represented by the COSO cube (COSO 2004, Exhibit 1.1, p 23). The front face of the cube represents the generic risk management process. The top face of the cube represents four generic business functions applicable to any organisation, and under which the objectives of the organisation may be categorised (Strategic, Operations, Reporting, and Compliance). The side face of the cube represents the hierarchical division of the enterprise into its various operating units. The cube thus represents the application of the risk management process to and across the functional areas of the business, and its integration throughout the enterprise, top to bottom.

### **Cross-functional integration – the portfolio view**

The first distinguishing feature of ERM is the notion of an holistic, cross-functional approach, in contrast to the traditional treatment of risks in functional “silos” (Lam 2003; Ward 2005; Power 2007; Ward 2006). Nocco and Stulz put it succinctly: “A corporation can manage risks in one of two fundamentally different ways: (1) one risk at a time, on a largely compartmentalized and decentralized basis; or (2) all risks viewed together within a coordinated and strategic framework” (2006, p 8).

Conceptually, the idea is simple. Organisations invest capital to achieve their objectives and recoup a return. At any given time, an organisation is exposed to multiple sources of risk (i.e. future issues, events, or scenarios which may potentially have an impact on the achievement of the company's objectives), some being unavoidable or inherent in the nature of the firm's operations and environment, and some arising from deliberate business decisions. Traditional risk management treats these sources of risk independently. That is, the identification, analysis, evaluation, and management of each source of risk is carried out largely independent of the same processes applied to other sources of risk. However, while some risks may be unrelated, others may not. Risks, and the treatment strategies employed to control one set of risks, may compound (exacerbate) or offset (mitigate) other risks. Such interdependencies mean that the “siloed” management of risks may (or will likely) result in sub-optimal outcomes at the enterprise level: (1) risks considered acceptable when assessed independently may result in an unacceptable level of risk when considered together; (2) cost

inefficiencies due to the unnecessary duplication of management processes and risk controls; and (3) the sub-optimal allocation of investment capital. In order to avoid or minimise these outcomes, an enterprise-wide view of risk is required:

In essence, a company should be viewed as a portfolio of businesses, each with its own unique risk/return characteristics.... Since the ultimate goal of management is to maximize shareholder value, the overarching principle for enterprise-wide portfolio management should be to manage the business portfolio in the same way that a fund manager manages a stock portfolio. In other words, business portfolio managers should strive to understand the links between risk origination (e.g., business lines and trading units) and risk transfer (hedging and insurance) and make those investment decisions that position the overall enterprise portfolio at the classic “efficient frontier” of risk/return... *(excerpt from Lam 2003, p 83-84)*

The central concept here is that of the Enterprise Risk Profile. This is conceived, in abstract terms, as an aggregate representation of a company’s risk universe (Cummings 2008), accounting for all types of risk exposures within the company, across different business activities and risk types, and for the inter-relationships between those exposures. As an object, the enterprise risk profile symbolises the idea that the (known) universe of contingencies potentially affecting a firm’s future performance can be captured in some form of representational media, and meaningfully interpreted by decision-makers. It is thus a defining concept of Enterprise Risk Management because it represents the state of knowledge required to make globally efficient business decisions.

## **Embedding risk management – creating the risk culture**

The second distinguishing feature of ERM is the notion of embedding risk management throughout the organisation and its processes, as a culture, rather than as an additional “bolted on” bureaucratic function (Lam 2003; Layton and Fuchs 2007; Shenkir and Walker 2006; Ward 2005). Again, the idea is conceptually simple. Risk management is more effective and efficient when it is carried out as an integral part of organisational processes, rather than as an added extra that people have to accommodate on top of their normal work. In this sense, the ultimate vision painted by world-level frameworks such as COSO and ISO 31000 is of a risk-aware culture where people are aware of and understand risk in the context of the organisation’s objectives, where all decisions in the organisation, from strategy setting to operations, involve the explicit consideration of risk to an appropriate degree, and where risk management procedures and tools are employed as a matter of course in day-to-day activities (COSO 2004; ISO 2008; Shenkir and Walker 2006). The COSO framework document

calls this “building enterprise risk management into the fabric of operations” (2004, p 18).

## Centering the corporate objectives

Common to each of the three dimensions of ERM, above, is an explicit centering of the corporate objectives as the focal point for all encounters with risk (Hutter and Power 2005a; Ward 2005). Individuals encounter risks in their day-to-day activities, but whether and how they identify, assess, and respond to those risks depends very much on their specific work context, experiences, skills, and points of reference. The risk management process (identification, analysis, evaluation, and treatment) seeks to encourage individuals to be deliberately and explicitly reflective about risks and controls in their specific work contexts. Enterprise Risk Management seeks to impose a common point of reference for risk assessments, so that risks assessed in one part of the organisation may be comparable with risks assessed in other parts of the organisation. That common point of reference is the organisation’s strategic objectives, and other such objectives which flow “from the strategy, cascading to entity business units, divisions, and processes” (COSO 2004, p 18).

In principle, the centering of the organisation’s objectives is fundamental to the performance of the risk management process, to the idea of the enterprise risk profile, and to the notion of creating a risk-aware culture in the organisation. The degree of potential impact on the achievement of the organisation’s objectives provides a common definition and basis for measuring risk (i.e. for evaluating the significance of the consequences of identified risk events). This, in turn, makes possible the aggregate representation of potential future impacts on the achievement of the organisation’s objectives arising from *all* identified sources of risk across the enterprise. For the latter, the corporate objectives, translated and made portable in risk frameworks, serve as a point of reference for nurturing a common and consistent understanding of risk throughout the organisation.

## Normative models for benchmarking ERM

Capability Maturity (CM) Models provide standardised benchmarks against which to compare the organisational development of ERM capabilities:

A capability maturity model (CMM) is a simplified representation of an organisational discipline... that distils industry practices into a coherent, process-based framework. These models are constructed according to maturity levels, from learner to best practice, which are characterised by the extent to which the processes are defined, controlled and institutionalised.... Capability

models enable organisations to establish their current level of process maturity and identify the steps necessary to progress to a higher level, building on their strengths and improving on their weaknesses. They may be used for benchmarking purposes, enabling organisations to compare themselves against other companies in their sector or beyond. (*excerpt from MacGillivray et al. 2007a, p 88*)

Most CM models are based the original Carnegie-Mellon Software Engineering Institute (SEI) model developed for the U.S. military for the evaluation of software subcontractors (Capability Maturity Model in Wikipedia 2008). The SEI model was itself based on a Quality Management Maturity model developed in the 1970s (Capabiliy Maturity Model in Wikipedia 2008; Crosby 1979). The SEI model has been adapted to a range process areas in various industries, including system engineering, project management, risk management, and personnel management (Capability Maturity Model in Wikipedia 2008; Champlin 2004).

Various institutions have developed normative CM models for benchmarking the maturity of ERM within organisations - including: the RIMS *Risk Maturity Model for Enterprise Risk Management* (2006; OCEG 2007); the Institute of Management Accountants *ERM Maturity Model* (Shenkir and Walker 2007, p 24); the IACCM *Business Risk Management Maturity Model* (2003); the RMRDPC *Risk Management Maturity Level Development Model* (RMRDPC 2002); a model from the IBM Institute for Business Value (Petit 2006); and a model developed for the water utility industry at the University of Cranfield, United Kingdom (MacGillivray, Hamilton, Hrudey et al. 2006; MacGillivray et al. 2007a, 2007b). PriceWaterhouseCoopers (2007) used their own proprietary model to evaluate Watercare's risk management capabilities in their capacity as external auditors. Watercare also self-assessed its asset management capabilities, including risk management, as part of the annual Water Services Association of Australia asset management benchmarking programme, the results of which were published in Watercare's asset management plans (Watercare Services Ltd. 2007b; WSAA 2003). Table V.2 summarises and compares four of the models identified above, including the Cranfield and WSAA frameworks.

Not surprisingly, the structure and content of the above ERM capability models is relatively consistent. The models typically define five maturity levels in terms of a range of business process attributes, which can then be used to assess the company's risk management processes. The models also typically identify a number of specific risk management processes that an organisation should have. The maturity levels are defined by a number of qualitative descriptors for each process attribute, which are essentially qualitative statements of the process characteristics which should be found at each level of maturity (see Table V.3 for a sample of typical adjectives).



**Table V.2.** Comparison of selected Capability Maturity Models for Enterprise Risk Management (references in text on previous page)

	<b>CMM 1: RIMS/OCEG</b>	<b>CMM 2: RMRDPC</b>	<b>CMM 3: Cranfield</b>	<b>CMM 4: WSAA</b>
<b>Maturity Levels</b>	1. Ad Hoc 2. Initial 3. Repeatable 4. Managed 5. Leadership	1. Ad Hoc 2. Initial 3. Repeatable 4. Managed	1. Initial 2. Repeatable 3. Defined 4. Controlled 5. Optimised	1. Informal / None / Sparse / Rarely 2. Aware / Minimal / Limited / Occasionally 3. Formulated / Moderate / Moderate / Often 4. Advanced / Advanced / Predominant / Usually 5. Robust / Complete / Total / Always
<i>defined in terms of...</i>		Definition Culture Process Experience Application	Scope Integration Verification and validation Feedback and organisational learning Stakeholder engagement Competence Resources Documentation and reporting	Capability: Process Development Process Documentation  Execution: Process Coverage Process Frequency
<b>Process Attributes</b>	Adoption of ERM-based approach ERM process management Risk appetite management Root cause discipline			
<i>and applied to...</i>	Uncovering risks Performance management Business resiliency and sustainability	Policies and leadership Resources and training Plans and procedures Measures and status Oversight and quality assurance	Strategic risk planning Risk acceptance criteria Risk analysis Risk-based decision making Risk response Risk monitoring Integrating risk management Supply chain risk management Change risk management Education and training Risk knowledge management	1.4 Risk Management 1.4.1 Responsibility and accountability 1.4.2 Risk policy (context definition) 1.4.3 Identification and recording of risks 1.4.4 Risk determination and costing 1.4.5 Risk prioritisation 1.4.6 Implementation of risk mitigations 1.4.7 Monitoring and review 1.4.8 Data and information 1.4.9 Systems and tools
<b>Risk Management Processes</b>				

**Table V.3.** Typical adjectives used to describe risk management process maturity

<b>Immature / Novice</b>	<b>Intermediate Levels</b>	<b>Mature / Advanced</b>
Ad hoc	Explicit	Institutionalised (integrated)
Absence of formal process	Formal process	Proficient
Absence of structure	Consistency of application	Permeating culture (extensive)
Implicit treatment of risk	Specific risk mngmt tools	Pro-active
Lack of risk knowledge	Supported	Flexible, adaptive, innovative
Lack of risk awareness	Strategic risks addressed	Organisational learning
Reactive	Some monitoring and control	Appropriate / fit for purpose
Risk averse		Verified

Sources for Table V.3: (MacGillivray et al. 2007a, 2007b; OCEG 2007; RIMS 2006; RMRDPC 2002; WSA 2003, 2008)

In essence the CM models describe capability maturity on a scale ranging from virtually non-existent, to fully integrated and proficient. Initially Risk Management is either non-existent within the organisation, or performed only on an ad hoc basis. At intermediate levels Risk Management is a distinct, explicit, and formal process within the organisation, consistently applied. Risk Management processes reach “maturity” when they become integrated as a proficient capability throughout all business processes (i.e. are an integral part of “the way things are done”), and are dynamically adaptive to the needs of the organisation. At this level Risk Management permeates the culture of the organisation, is almost intuitively pro-active rather than reactive, is flexibly adaptive to suit a range of contingencies, and continuously evolves as organisational members learn from experience.

Capability Maturity models are, however, normative rather than descriptive. They purport to define the characteristics of effective enterprise risk management, but those definitions are not necessarily based on sound empirical analysis of actual risk management processes in organisations. There has, as yet, been no attempt to correlate enterprise performance with risk management process characteristics found in organisations (see Chapter 1). Rather, CM models for ERM appear to derive their process definitions and attributes from a top-down translation of the principles of ERM, with the various models achieving this in more or less detail.

## **Risk assessment in engineering contexts**

The core element of Enterprise Risk Management is the standardised Risk Management process, since it is this process which is to be applied across all business functions and

integrated throughout all business processes (COSO 2004). This section describes the phases of context definition and risk assessment, since it was with these phases of the risk management process that Watercare's Corporate Risk Manager was primarily concerned. The latter term "risk assessment" refers collectively to the phases of risk identification, risk analysis, and risk evaluation (as defined in AS/NZS 4360:2004). This section describes, at a conceptual level, what is involved in defining the system context, and identifying, analysing, and evaluating risk from a functional systems perspective; in other words, an engineering approach to risk assessment. This serves as the normative ideal against which the CRM's proposals for improving the quality of Watercare's "risk data" are compared in Chapter 3.

### **Why a functional systems perspective?**

Conceived as a process of inquiry, risk assessment seeks to answer several questions (Kaplan and Garrick 1981): (1) What are the outcomes (consequences) of interest? (2) In what ways might those outcomes be realised? And, since the future is uncertain, (3) How likely is it that those outcomes will actually occur within the given time frame? The answers to these questions are often represented as a statement linking some perceived source of risk to anticipated outcomes via some chain of causal events. However, what individuals and societies perceive as risk is shaped not just by the objective state of risk, but also by social, cultural, political, and technological factors (Rosa 1998, p 28; see also Hilgartner 1992)<sup>28</sup>. It is therefore possible that any state of the world implying an intersection of human stakes and uncertainty may be legitimately conceptualised as risk. In this sense, the normative scope of the label of risk is virtually unlimited, potentially encompassing any domain of phenomena in the world. Depending on the systems (context) definition adopted, whether implicit or explicit, risk may refer to physical, chemical, biological, political, cultural, or organisational phenomena, or all of these in combination.

World-level standards for ERM present a generalised model of the risk management process. That process may be applied in very wide range of contexts (different types of organisation, different industries, different functional contexts). Consequently, the standardised risk management model assumes no particular epistemological orientation with respect to the task of knowing risks, and world-level frameworks and guidelines provide very little methodological guidance for risk assessments. Rather, decisions about

---

<sup>28</sup> See also discussion in Chapter 4.

phenomenological foci, epistemological orientations, and methodological traditions are left up to those performing the risk assessment; since the methodology of the inquiry necessarily depends on the nature of the phenomenon being investigated, the questions being asked, and who is asking them. Even if risk assessments everywhere ask common questions (Kaplan and Garrick 1981), very different disciplinary conceptualisations of risk and approaches to risk analysis may be adopted depending on the context of the assessment and who is performing it (Althaus 2005).

While this epistemological variability goes some way toward explaining why world-level risk management standards and frameworks such as AS/NZS 4360 do not provide specific guidance on how to perform risk assessments, it also poses something of a problem with respect to defining a normative ideal against which to compare the Corporate Risk Manager's vision and proposals for risk assessment at Watercare. Since organisations are exposed to myriad types of risk, this variability must (or should) necessarily be reflected in the approaches to risk analysis employed within the various operational contexts of the organisation. Thus, for the purpose here, which tradition of inquiry should serve as the normative basis for comparison?

In answering this question I chose to draw on the observed fact that the dominant notion of risk within the cultural milieu of the Watercare organisation was that associated with the potential failure of physical assets within the water and wastewater infrastructure networks (see Appendix I). While this may be seen as a form of bias from an ERM perspective, which demands an integrative view of risk encompassing all of the risk 'silos' across the enterprise, it was not unreasonable in light of the fact that developing, operating, and maintaining infrastructure assets was Watercare's core business. As such, the company's primary business activities (operations, maintenance, planning, project delivery) were engineering activities, and the bulk of the company's employees were engineers, engineering technicians, or scientists; this included the CRM, who held a PhD in fire engineering. Thus, at Watercare, engineering concepts of risk and engineering approaches to risk assessment predominated. On this basis, the following section explicates an approach to risk identification and assessment based on the functionalist systems methodology common in engineering (Auyang 2004; Ayyub 2003; Jackson 2000).

## **Engineering ideals for risk assessment**

Engineering is "the profession in which knowledge of mathematical and natural sciences gained by study, experience, and practice is applied with judgement to develop ways to

utilise, economically, the materials and forces of nature for the benefit of mankind” (1982 definition from the Accreditation Board for Engineering & Technology). Technology is the outcome of the engineering process – engineers engage in creative processes of design which seek ways to manipulate certain features of the natural and physical world to produce technologies which fulfil human objectives (National Academy of Engineering 2004). Although engineering is very much a social activity requiring interaction with a great many people, the principal domain in which engineers seek to intervene is the natural and physical world, and the principal knowledges that engineers seek to apply have their foundations in the natural, physical, and mathematical sciences.

For engineers, risk arises from the behaviour and interactions of very real systems and processes, which must be understood through scientific analysis. Inquiry in engineering predominantly adopts a functionalist systems methodology (Auyang 2004; Ayyub 2003; Jackson 2000; Keey 2000; see also Appendix II). A claim to a functionalist systems methodology thus implies an assumption that the real world is systematic, an analysis of the problem conducted in systems terms, including the construction of models aiming to capture and represent the real world system, and a presumption that quantitative analysis is useful (Jackson 2000, p 203). Risk assessment in engineering makes this claim explicitly:

Risk studies require the use of analytical methods at the system level that takes into consideration subsystems and components when assessing their failure probabilities and consequences. Systematic, quantitative, qualitative, or semiquantitative approaches for assessing failure probabilities and consequences of engineering systems are used for this purpose. A systematic approach allows an analyst to evaluate expediently and easily complex systems for safety and risk under different operational and extreme conditions.... Risk assessment is a technical and scientific process by which the risks of a given situation for a system are modeled and quantified. (*excerpt from Ayyub 2003, Risk Assessment Definition, p 43*)

The following sections explicate the stages of context definition and risk assessment using a functionalist systems perspective. The latter stages of risk treatment and monitoring are not dealt with here since the CRM’s primary focus during the research period was on context definition and risk assessment, and these are consequently where the attention of this thesis is directed. The intention in the following sections is to describe at a conceptual level what is involved in making and evaluating statements of risk, reflecting the notions of risk and inquiry which are dominant in the discipline of engineering (see, for example, Ayyub 2003; Hall et al. 2006; Keey 2000; Le Masurier, Blockley, and Wood 2006; Lewin 2006; MacGillivray, Hamilton, Strutt et al. 2006).

## Defining context

To identify risk is to establish the possibility of a causal connection(s) between a certain source(s) of risk and a certain outcome(s). Risk identification may proceed either forward or backward (i.e. inductively or deductively, Ayyub 2003). That is, one may start with certain outcomes of concern and, by asking “how might those outcomes arise?”, work backwards to identify the sources of risk; or one may start with potential sources of risk and work forwards to determine the resultant outcomes (which may or may not be of concern). Both procedures involve the bounding and definition of systems (processes); in other words, the definition of context, which is the first phase in the generic risk management process.

A system may be defined as a group of components (which may be systems themselves) connected in a configuration that allows the system to perform specific functions within a defined boundary across which inputs and outputs flow. The system relies on inputs of materials and energy from its environment in order to function, and produces outputs as a result of that functioning. The term “system” strictly refers to the group of components and their relationships (rules of interaction). It has no temporal component and is therefore a static concept. Systems are physical entities which simply exist and do not change unless acted upon by a process. They are brought into existence by processes (i.e. development), their existence may give rise to processes (i.e. system functioning), and their existence is terminated by processes (i.e. decay, destruction, change/evolution). The term “process” means “a series of actions or operations conducing to an end”. It refers to the actual functioning of a system, including all of the actions and interactions of the system components, the inputs and their transformations, and the outputs produced; ‘process’ is therefore a dynamic concept. Further, processes transcend the boundaries of systems (i.e. a system may perform functions within several processes) such that, from a process perspective, the boundaries of a system are meaningless unless they also constitute the boundaries of the process(es) of which that system is a part. In the sense used here, the terms “system” and “process” are not limited only to the physical world (e.g. buildings, machines), but also encompass what is commonly referred to as the social (e.g., people, teams, relationships). Drawing on Latour (2007), the concepts of “system” and “process” should be understood as defined by the connections (relationships) between things, rather than the things themselves.

Context sets the scene for identifying and defining risk. Defining context is the act of putting in place boundaries around the relevant systems (processes) which give rise to sources of risk, and which are affected by the resultant outcomes. The assumption is that

everything that humans value is the product of a system (process). Whether it is human life and health, love and friendship, development, nature, or financial return, all are the products of multiple interacting systems (processes), constituted by physical, ecological, and technological components, as well as by people and their actions. In such a world, everyone and everything is part of multiple systems (processes), and each system (process) is linked to multiple others by transfers of energy (i.e. inputs and outputs; energy may take a variety of forms including material, electrical, chemical, or physical force, or information). Changes in one system (process) thus affect the functioning and performance of multiple other systems (processes) through the resultant changes in those energy transfers. These changes thus affect the production of things that humans value, sometimes positively, sometimes negatively, sometimes significantly, sometimes negligibly.

A company, for example, produces value for its shareholders, the consumers of its products and services, and its employees and other stakeholders. The company produces that value through the functioning of various systems (e.g. business units, physical systems, assets, etc.) and processes (e.g. planning processes, reporting processes, performance appraisal processes etc.). Its performance over time, and the sustainability of the value that it produces, is a function of the ongoing performance of those systems and processes in interaction with factors in the external environment (e.g. competitors, regulation, natural resources). The company's shareholders value the financial returns on their capital, and, for them, outcomes of concern will be events which affect the company's revenue and profits, either positively or negatively. The sources of these outcomes will be either lapses or improvements in the performance of the company's business systems (processes), or external events which affect the company, but over which the company might have little or no control.

Risk is, in this sense, uncertainty about the future variability in the performance of the organisation arising from potential changes in the internal and external environments; which is precisely the definition used by Ward (2005, p 45): “[risk is] the implications of significant uncertainty about the level of performance achievable... we can then associate ‘downside’ risk with the implications of significant ‘threats’, or unwelcome consequences, and ‘upside risk’ with the implications of significant ‘opportunities’ or welcome consequences.”

For the purposes of assessing risk, a system or process is relevant to the extent that the functions it performs and the outputs that it produces contribute to, or otherwise impact upon, the achievement of a specific goal or objective of value to humans. There are, thus, two main parts to defining a system (process): (1) requirements analysis – defining the specific goal or objective, and (2) functional analysis – defining the system boundaries, components and measurable performance criteria, and collecting information for assessing failure

likelihood and consequences (Ayyub 2003, p 45). In other words, the context for risk identification is established by defining, first, what we value (the definition of 'ends'), and, second, the systems (processes) that produce what we value (the definition of 'means'). The specification of 'means' begins by defining a boundary around the system (process) of interest. The consequence of this distinction is the simultaneous definition of both what is internal to the system (process) in question, and what is external. Since no real system (process) is completely isolated, the system definition can not be complete until what is outside the boundary has been related to what is inside in terms of what crosses the boundary (i.e. by specifying input relationships). These distinctions thus define the system (process) of interest for the analysis, which may proceed once the relevant time frame has been specified, and the appropriate assumptions regarding the state of the environment over that time frame have been made. Table V.4 summarises these contextual elements of risk.

*Table V.4. Contextual elements of risk*

<b>Element</b>	<b>Description</b>
Definition of objectives	This is the definition of 'ends'. Risk is defined as a future situation in which something of human value stands to be lost or gained, but where the realisation of that loss or gain is uncertain. The definition of risk must therefore begin with the specification of that which is valued, and which stands to be gained or lost. Practically, for organisations (including projects), this means defining the performance objectives to be achieved.
Definition of system (process)	This is the definition of the 'means' by which what we value is produced, i.e. the specification of the systemic structures and arrangements which constitute the system (process) in question. Risk events may then be defined in terms of changes in the behaviour of system (process) components (e.g. failures) which affect the overall performance of the system (process).
Definition of the system relationship with environment	System (process) performance depends not only on the integrity of the internal components, but also on the relationships between those components and the surrounding environment. Input relationships must therefore also be defined.
Definition of time period	A statement of risk constitutes a prediction about the future, so it is necessary to state the time period over which that prediction applies.
Definition of environmental assumptions	Any prediction about the future must include a specification of the underlying assumptions about what the future will hold. That is, although in principle anything is possible, the objectives and system (process) already defined are predicated on certain assumptions of the nature of the future. Typical assumptions include future environmental conditions, the continuity and consistency of system (process) inputs, the nature of feedback loops, possible external influences etc. Without knowledge of the underlying assumptions it is impossible to meaningfully understand the results of the risk analysis.



## Identifying risk

The making of the above distinctions constitutes the construction of a model which represents the system behaviour of interest (Ayyub 2003). There are a wide variety of established analytical methods available for defining system models; for example, work breakdown structures, contributing factor diagrams, decision-analysis method, bayesian networks, process modeling method, black-box method, state-based method, component-integration method (Ayyub 2003, Chapter 3). With at least a basic model defined, risk identification may proceed by identifying how the achievement of the specific objectives might be compromised; in essence by asking “what could go wrong?” (Ayyub 2003).

At a conceptual level, there are three general ways in which things can “go wrong”, constituting the three general sources of risk for any system or process: (1) failure of an internal system or process component, (2) breakdown of an input relationship, or (3) deleterious impacts on the system or process from some external event. I am referring here to down-side risk, but the same three categories apply, conceptually, to up-side risk, or opportunity, simply by replacing the negative adjectives with positive ones. Within the systems definition, these three categories are necessary and sufficient to describe any and all possible risk events for that system. They are necessary because they represent the three fundamental ways in which any system or process may fail to achieve its objectives, and they are sufficient because they may be combined in any number and order to represent a causal chain of events of any length or complexity.

On this basis, and ignoring probability for a moment, a statement of risk consists of some combination of the above three categories (i.e. the chain of causality) linked to a fourth category, that of the anticipated consequences of the risk event, defined in terms of impact on the system objectives. Figure V.1, below, represents this conceptual relational structure for risk statements. The process of risk identification entails identifying potential hazards, events, and scenarios, and making them explicit as risk statements of this form in Figure V.1. Again, there are a wide variety of methods available for identifying risks, ranging from informal (e.g. brainstorming, what-if analysis) to highly systematic methods such as Failure Mode and Effects Analysis (Ayyub 2003, Chapter 2; MacGillivray, Hamilton, Strutt et al. 2006).

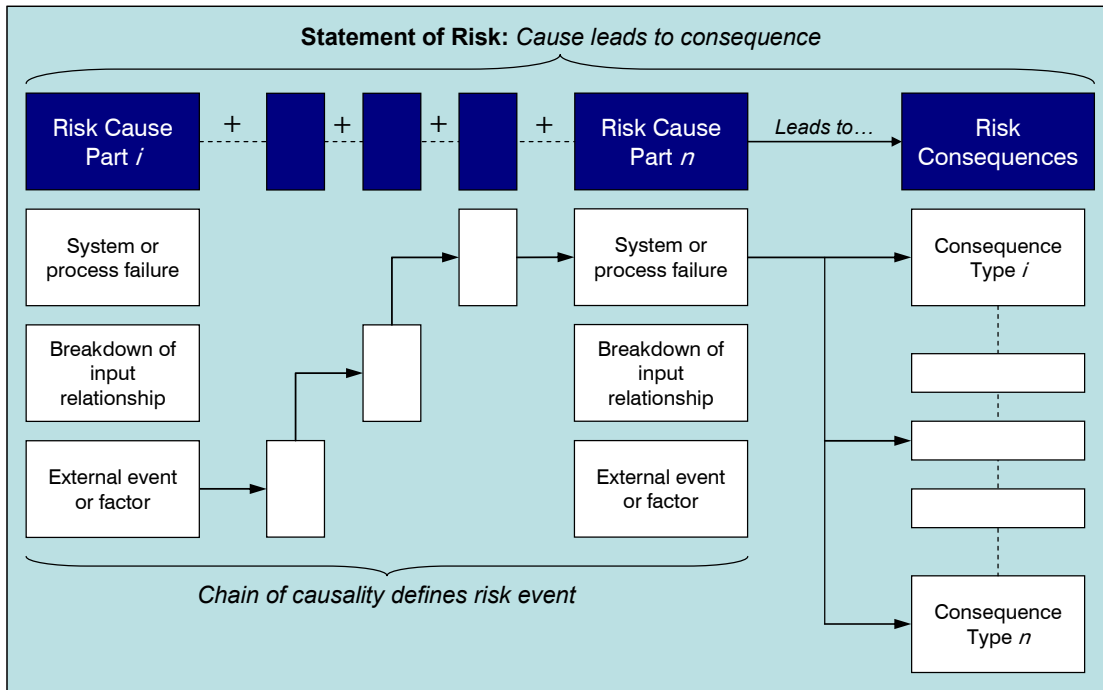


Figure V.1. Generic relational structure of statements of risk

## Analysing risk

Once the possibility of causal connections between risk causes and their consequences have been established, risk analysis seeks to assess the magnitude of the resulting statement; that is, to calculate measures of the magnitude of the anticipated consequences and the probability that the specified consequences will eventuate, within a specified time frame, arising from the specified source and chain of events. However, while risk analysis is often portrayed as a quantitative scientific process, particularly in engineering and economics (Althaus 2005, Ayyub 2003, Keey 2000), the precise quantification of risk is not often simple. The linear representation between cause and consequence in Figure V.1 is a useful pedagogical tool for illustrating the concept of a statement of risk, but the complexity of most real world systems means that the relationship between cause and consequence is rarely, if ever, only linear and one-to-one. There may be multiple ways in which a single outcome could occur, or a single source of risk may give rise to multiple outcomes. Further, in most real world situations, it is unlikely, given the occurrence of a particular originating event, that every subsequent event in the causal chain will be 100% certain to occur. A clear example of this is where layers of redundancy in a network or system accommodate for failures. Indeed, some systems may be sufficiently complex that it may be impossible to make a clear statement linking cause to consequence; leaving only the possibility of calculating the

statistical probability of certain outcomes based on historical system behaviour at a macro level.

Realistically then, any statement of risk implies a distribution of possible outcomes, some of which will be more likely than others. The mathematically correct quantification of such statements requires the use of probabilistic methods (Wood 2008). However, the ability to accurately model systems and calculate and represent risk depends upon a number of factors, including the nature of the system (process) in question, the available knowledge about that system (process), and the calculative methods and technologies available (Ayyub 2003; Keey 2000; HB 436:2004)<sup>29</sup>. The level of detail in the analysis is also usually constrained by the purpose of the assessment. In practice, a variety of methods may be used for ranking risks; from simple qualitative two-by-two matrices, to more sophisticated semi-quantitative matrices (such as the likelihood and consequence matrices employed within Watercare's risk framework), to fully quantitative probabilistic methods (Ayyub 2003; Keey 2000; HB 436:2004). The following section describes Reliability Centred Maintenance modelling as perhaps the ideal standard for quantitative risk analysis. Fully quantitative methods are, however, resource, time, and data intensive, and their use is generally limited to situations where the perceived level of risk is significant, where the contextual complexity is high, or where mathematical precision is required (e.g. research). Qualitative and semi-quantitative methods, which tend to forgo mathematical accuracy in favour of simplicity, employing discrete measures of consequence and probability, are often used where data, time, and resources are limited, where the perceived level of risk exposure is low, or where contextual complexity, and hence uncertainty, is low.

The resulting combination of cause, consequence, and probability constitutes the final statement of risk, the magnitude of which is represented by the measures of the latter two parameters in combination (consequence × probability). The full statement of risk, incorporating a description of the causal event, is necessary for articulating risk controls (i.e. risk controls may be targeted to eliminate the cause, or reduce the probability or consequences). The magnitude of the risk is commonly represented via some form of bi-axial matrix or chart, such as those used within Watercare. The intention of such representations is to rank the identified risks from most significant to least significant. With the list of risk statements so prioritised, it then becomes possible to make decisions about which risks require treatment and the level of resources to commit to those treatments (HB 436: 2004).

---

<sup>29</sup> HB 436:2004 is the Companion Guidelines Handbook to AS/NZS 4360:2004.

## **Reliability Centred Maintenance modelling (the quantitative ideal)**

Reliability Centred Maintenance (RCM) is a rigorous, comprehensive, rational methodology for optimising preventative maintenance schedules for engineered systems. The primary objective of RCM is to preserve system function (Smith 1993). The methodology, of which there are many variations, employs a whole-of-system approach to identify modes by which system function may be compromised as a result of component failure (formally referred to as Failure Mode and Effects Analysis, or FMEA), and to quantitatively model the reliability of the system function using available data on the reliability of the individual components (Ayyub 2003; Moubray 1997; Smith 1993). The analysis is used to identify and prioritise those components which require condition monitoring and preventative maintenance to minimise the likelihood of failure, and those components for which no proactive intervention is required (i.e. they can be run to failure).

The reliability of components (products, devices, etc.) designed and produced by the same production process is variable due to uncertainties associated with materials, the manufacturing process, and the effects of different operating environments (Ayyub 2003; Moubray 1997; Smith 1993). The basis for the RCM methodology is the observed fact that despite these uncertainties the global rate of failure for any population of components of the same type follows a discernable and predictable pattern over time. One version of this pattern is that the rate of failure is constant over most of the equipment lifespan, but increases exponentially toward the end (referred to as the wearout period). Another version of the pattern includes an “infant mortality” period where the rate of failure within the population is initially high but decreases to a constant rate. There are six typical failure distributions, but all feature a dominant normal life period where the rate of failure is low and reasonably constant, i.e. where equipment failures are randomly distributed (Ayyub 2003; Moubray 1997; Smith 1993).

The reliability (time to failure) of any given component or product can be treated as a random variable and probabilistically modelled, and the failure rate distribution during different life phases can be modelled using a Weibull distribution (Ayyub 2003; Moubray 1997; Smith 1993). For a specific component, the probability density function (pdf) of the time to failure may be available as a global statistic, obtainable from the manufacturer or via some industry database (e.g. the pdf for a specific type of lightbulb). For other components, however, global statistics may not be applicable due to the uniqueness of particular operating conditions. In this latter case, the pdf of the time to failure may be estimated from the maintenance and repair logs for the particular component in question. Such logs provide in-

use life data, typically consisting of recorded times to failure (for non-repairable components), times between failures (for repairable components), and, in some cases, recorded time intervals not terminating in failure (Ayyub 2003; Moubray 1997; Smith 1993). There are also established methods for incorporating more subjective prior information (e.g. expert opinions, experience) using Bayesian inference.

With the reliability of individual components represented by known probability density functions, it subsequently becomes possible to quantitatively model the reliability of an entire system via Monte Carlo simulation (Ayyub 2003). Reliability models calculate a probability density function for the reliability of the system as whole by simulating, over several thousand iterations, the performance of the system components over a set period of time (Monte Carlo simulation). For each iteration the simulated failures of individual components are randomly distributed according to their own respective probability density functions. The output of such availability simulations is a quantitative estimate of how often the system function will be compromised by component failure within a specified time period (Ayyub 2003). The results of the simulation can be interrogated to determine which components are most often expected to cause system down-time (i.e. the weakest links). Proactive measures can then be taken to improve the reliability of those components, such as redesign, installing redundancy, or rescheduling preventative maintenance (Ayyub 2003; Moubray 1997; Smith 1993). For highly complex processes, involving multiple parallel and series systems with thousands of components and many layers of redundancy, availability simulation is the only way of analysing system behaviour to identify critical components.



# Appendix VI

## Experiences with implementing ERM

The sections of this appendix describe the CRM's experiences with the following tasks:

- Framework development: Redefining Watercare's existing one-size-fits-all risk assessment framework into three separate but related frameworks for the assessment of risk in the Strategic, Operations, and Project Management contexts of the business.
- Tool development: the CRM redesigned the Project Risk Register to capture "risk data" in a format consistent with the redefined Project risk assessment framework.
- Modelling: Promoting the development probabilistic risk models (Monte Carlo simulation) to quantify the Enterprise Risk Profile.
- Facilitation: the CRM also engaged in the direct facilitation of risk identification and assessment workshops with staff and management.

### **(Re)defining Watercare's risk assessment framework**

The CRM perceived that many of the problems with the quality of Watercare's "risk data" stemmed from problems with the company's risk assessment framework (described in Chapter 3). The redefinition of that framework was therefore one of the first tasks which the CRM undertook and one to which he devoted considerable attention. In this Appendix the term "framework" refers explicitly and only to the physical document which specified the

categories and evaluation criteria by which risk was to be assessed within the company (at Watercare this was a table with five consequence categories and six levels of severity). This is a rather narrow usage of the term “framework” in contrast to its broader meaning which also encompasses policies, rules, and procedures for risk management (see Appendix 1 for a broader description of Watercare’s risk management framework)

## **The design task**

The CRM wanted to reformulate the risk framework as an explicit deconstruction of the corporate objectives, such that the resulting risk hierarchy would more accurately represent the various operational contexts of the business (see the description of the CRM’s vision in Chapter 3). In designing a new framework the CRM had to address three questions: (1) How many categories should the framework have and what should those categories be? (2) Which criteria should be used to define risk consequences in each category? And (3) How should those criteria be ranked to denote the relative significance of outcomes, both within and across categories?

The CRM anticipated that the first question could be addressed by deconstructing the corporate objectives in Watercare’s Statement of Corporate Intent, i.e. by mapping the corporate objectives in that document to certain categories and sub-categories within the risk framework. Rather than a one-size-fits-all framework, the CRM envisaged a system of three inter-related frameworks which, when fully developed, would specify how risk was to be evaluated within the company, each framework corresponding to a different management context (Clement 2007d). The Operations Risk Framework would specify criteria for evaluating risks in the context of the company’s day-to-day activities delivering water and wastewater services, the Project Risk Framework would specify criteria for evaluating risks in the with respect to developing and delivering changes in performance reliability and capability, while the Strategic Risk Framework was intended for use within the longer-term business planning context (Clement 2007d).

Three frameworks were required because the relevant objectives, time scales, and magnitudes of the risks involved in each context were quite different. The Operations and Strategic risk frameworks would be used to define, justify, and prioritise business development projects and were therefore intended to measure impacts against the fundamental objectives of the business (see Figure VI.1). The Project risk framework would be used to evaluate risks in the context of delivering those projects and was therefore intended to measure impacts against project specific deliverables (see Figure VI.1).



<p><b>Risk assessment categories in the Business Operations Risk Framework (October 2007):</b></p> <ul style="list-style-type: none"> <li>▪ Delivering water</li> <li>▪ Treating water</li> <li>▪ Being economical</li> <li>▪ Maintaining statutory compliance (includes wastewater objectives)</li> <li>▪ Maintaining stakeholder confidence</li> </ul>	<p><b>Risk assessment categories in the Business Projects Risk Framework (October 2007):</b></p> <ul style="list-style-type: none"> <li>▪ Delivering the project to schedule</li> <li>▪ Delivering the project within budget</li> <li>▪ Delivering an outcome that satisfies the business need</li> <li>▪ Maintaining statutory compliance during the project</li> <li>▪ Minimising the impact of the project on internal and external stakeholders</li> </ul>
--	--

*Figure VI.1. Different objectives for different contexts  
(source: Clement 2007d)*

In order to support his vision of generating quantitative enterprise risk profiles, the CRM wanted the framework to contain objective performance measures: “[t]he thing that stands out immediately is that if you want to quantify [those risks] then you need some measure of performance, so some measure of service delivery standards or requirements, and some measure of business process or management requirements” (Dialogue 8, p 299). The framework would also need to be comprehensive, “[i]t’s about having an objective way of measuring things across an organisation, full breadth and full depth.” (Dialogue 3, p 257).

Rather than attempt to arbitrarily define the criteria by which the successful achievement of the corporate objectives should be judged, the CRM explicitly looked for and appropriated existing standards and criteria. In particular, he identified that there were already a number of external standards to which the company was required to respond (statutory obligations and the customer contracts), and that the company’s performance was also regulated by a number of internal standards (the operational budget, the Asset Management Plan, and the Funding Plan – these being the outputs of Watercare’s annual business planning cycle; see Watercare Services Ltd. 2006, Fig. 21, p 25). The CRM expected that the consequence evaluation criteria in each category of the new framework could be drawn from the relevant performance standards with which the company complied. So, for instance, one of the relevant standards for potable water quality was the New Zealand Drinking Water Standards (NZDWS), which specified a range of water quality related criteria. The CRM anticipated that it would be possible to collate or otherwise rationalise those criteria to produce appropriate indicators to represent risk impacts.

To be effective, the risk framework would have to balance breadth and detail against ease of use. If the framework was too complex and detailed in terms of the number of evaluation criteria and risk measures then it would be difficult to use. On the other hand, if the

framework was too simple and generic then it would neither address the problems of subjectivity and ambiguity which plagued the existing (2003) framework, nor provide for the calculation of quantitative enterprise risk profiles. As I commented in Dialogue 8, the question was “how do you group those commitments, and then how do you represent each of those categories? I mean, the measures you choose will effectively represent those categories, so represent risk, and what you’re trying to avoid is a situation where the measures are ambiguous, which would cause confusion about how to evaluate risks” (p 299).

## **Engaging with stakeholders**

The CRM’s approach to redefining the risk framework led him to engage with key stakeholder groups within the company; specifically:

- The CRM met with project managers from the Asset Management Group in June 2007 in order to define the functional context in which they worked, and to present his initial proposal for a revised Project Risk framework.
- The CRM met with planning and project managers from the Asset Management group in July 2007 to further investigate the question of how project risks impacted the enterprise.
- The CRM met with line managers from the Operations group in August 2007 to discuss risk assessment parameters and severity scales for the proposed Operations Risk framework;
- The CRM met with a group of Watercare’s general managers in October 2007 to present his revised Operations Risk framework.
- The CRM engaged with Watercare’s general managers and senior line managers, over a period of several months in late 2007 and early 2008 for the purpose of defining the Strategic Risk framework.

The first four of the above engagements are described in the following sections. The latter evolved into a fully fledged facilitative process and is described later.

### **Engaging with stakeholders: Planning and Project Managers**

The CRM met with project managers from the Asset Management Group in June 2007 (see Dialogue 5) and with senior planning and project managers in July 2007 (see Dialogue 6). The

purpose of the meetings was for the CRM to understand the functional contexts of the planning and project teams, and to address the question of what were the enterprise-level impacts of planning and project risks (i.e. “how does Watercare feel the impacts of these risks?”). In the June meeting the CRM sought to elicit information about the processes and work flows within the Projects team, the relationships between Projects, Planning, and Operations, the level of engagement between the Projects team and external stakeholders (e.g. the Local Network Operators), and what processes and practices the project managers employed to manage risk. The July meeting focussed primarily on the latter question of how the broader Watercare enterprise actually “felt” the impacts of “cock-ups” by the Planning and Projects teams. Several points were notable.

First, the clearly defined administrative boundaries of the organisation obscured a more fluid processual relationship between the three main functional groups in the organisation. In particular, the project managers complained that the hand-over point between Planning and Projects had been gradually shifting over the years such that the Projects team often had to undertake planning and design work that they felt should be the responsibility of the Planning team. As one project manager commented: “...really the investigation should be thorough enough to give us all the information that’s necessary for a Capex request. But what we’re getting is back-of-fag-packet jobs, straight lines on a drawing, “there you are guys, that’s what you need to do, get on with it”” (Dialogue 5, p 275). This was notable from a risk management perspective because the hand-over of a project from the Planning team to the Projects team also corresponded to a transfer of responsibility for managing the risks on that project. If the Planning team did not adequately scope the risks during the initial planning phases, the Projects team could inherit risks over which they might have little control, and for which they did not have the resources to properly manage. Indeed, it had been a common gripe amongst project managers (and also some Operations staff) that they were getting blamed for project delays and cost-overruns where they felt the root cause was poor initial scoping of the project by the planning team.

Second, while there was potential for a “cock-up” on a project to cause a significant immediate impact in terms of financial cost or service delivery failures, for the most part, the impacts of project management failures were considered to be relatively insignificant from an enterprise-wide perspective, or were not “felt” immediately. For instance, as one project manager commented, “if there’s a bit of a collapse and you lose a bit of pipe and it costs a

\$100,000, what's that in the big scheme of things?"<sup>30</sup> (Dialogue 5, p 281), while another manager commented that "[g]enerally it wouldn't matter if [a project] was two or three years late. And it shouldn't matter... [s]omething has gone wrong at the planning level if we need to suddenly have a project delivered tomorrow" (Dialogue 6, p 284). The most significant impacts from planning and project management failures were actually the potential long-term consequences:

- If an asset was not delivered to specification, the effects would manifest as operability problems and increased operating and maintenance costs over the life of the asset.
- If a project was delayed for too long, not delivered at all, or failed completely, then the infrastructure would not be in place for Watercare to meet its future obligations. The effects of this would manifest as a gradual degradation in current levels of service accompanied by a gradual increase in risk exposure as operational head-room decreased, or as a failure to meet some new level of service required at a certain point in the future (e.g. meeting new water quality standards or resource consent conditions).
- Poor asset management planning on a systems level would also manifest over time as a gradual degradation in levels of service and increased risk exposure, or as an increasingly "peaky" work flow accompanied by significant future increases in capital expenditure (and, consequently, increases in water and wastewater prices). This was in fact the situation that the company faced in 2007. The company had delivered an Asset Management Plan at the end of 2006 which signalled large Capex and price increases over 5 – 10 years to address several prior years of under investment on the Water side of the business.
- Persistent planning and project management failures could also adversely affect the organisation's reputation, particularly with customers (the Local Network Operators), shareholders (the local authorities of the Auckland region), and with business partners (e.g. contractors and suppliers).

The recognition of these potentially significant long term impacts highlighted how the Asset Management group, even more so than Operations, was fundamentally responsible for

---

<sup>30</sup> A project budget overrun of \$100,000 would amount to 0.1% of Watercare's total annual average Capex budget \$NZ100 million.

the company's long term performance in terms of its primary legislative objective to balance the long-term costs of delivering water and wastewater services against the long-term risks to the delivery of those services. In this regard, while project-level risks were, individually, often of relatively minor significance at the enterprise level, the systemic, enterprise-wide risk posed by consistent under-performance of the Planning and Project Management units was considerable.

This situation was problematic with respect to the assessment of project risks. The fact that the immediate consequences of project risks, such as schedule and budget overruns, were often individually insignificant when viewed in the context of the enterprise as a whole, even if they were significant as a percentage of the original time frame and budget assigned to a particular project, meant that Watercare's one-size-fits-all risk framework was inappropriate for assessing project-level risks. Assessing project risks with the same severity scales as were applied to outcomes in the Operations context would always lead to one of two outcomes. Either the project risks would always be rated as negligible (i.e. Class 1 or 2), or the project managers would have to fudge the assessment to make sure that the resulting risk score matched their perception of the magnitude of the risk. Indeed, the project managers openly admitted that the risk scoring process was "messy", and that they often had to manipulate the numbers: "... when you have the feeling that certain risk needs to be a Class 4 cos it is important so you make the numbers right that it gets a Class 4. So the numbers don't really help a lot. The classification is useful, but tweaking the numbers to get the right classification doesn't really help" (Dialogue 5, p 279).

For the CRM, looking at the problem from the perspective of designing an integrated risk assessment framework, the issue was one of how to give project risks a score which appropriately reflected their significance, without being inconsistent with the scoring system used for Operations risks. The CRM's solution, which he worked out some time later, was to define significance at the project level in terms of the responsibility and accountability of individual project managers. Significance at the project level was a function not of the degree of impact on the Watercare enterprise, but of the degree to which the impact would reflect on the competence of the individual project managers. Thus, risks which would otherwise be considered negligible in the enterprise-wide context could be rated as significant when considered in the context of the individual project managers' performances (i.e. a cost overrun of \$50,000 on a project with an original budget of \$50,000 reflects poorly on the project manager, even though the amount is negligible in the context a total enterprise capital spend of \$100 million). Under the CRM's framework, project risks were therefore reported with two scores: the Project-level risk score, as just described, and the Corporate-

level risk score, adjusted to reflect the significance of the risk from the enterprise-wide perspective. This solution also addressed, to some extent, the aforementioned problem of the potential long-term significance of Watercare's cumulative risk exposure across all projects by making the performance of individual project managers a key risk control. In this regard, Dialogue 5 revealed that, in addition to the project risk register, Watercare's project managers already used a range of tools and practices to effectively manage project risks, including: a "shut-down register" to co-ordinate asset shut-downs with the Operations group; getting "buy-in" from the Operations group during the detailed design phase and maintaining that dialogue during the construction phase to ensure that the delivered asset would meet all operational requirements; and establishing compliance registers to manage compliance with consent conditions from external agencies (e.g. roading authorities, the Auckland Regional Council).

The information acquired by the CRM from these meetings with Watercare's planning and project managers informed the design of his revised project risk framework. The details of that framework, and a description of its use on a major project, are presented later.

### **Engaging with stakeholders: Operations Managers**

The CRM met with several line managers from Watercare's Operations group in August 2007 (see Dialogue 4). The CRM had called that meeting early in the development of his revised risk framework to address the question of how to represent the corporate objectives with concrete parameters that could be quantified using common business data. For the CRM it was important that his revised framework should use parameters that Operations staff were familiar with on a day-to-day basis. To this end, the CRM had explicitly looked for and appropriated existing standards and criteria which applied to the company's operations. In the August meeting with the Operations line managers the CRM presented a draft framework which described risk outcomes ranging from "Negligible" to "Catastrophic" across five evaluation categories: impacts on water delivery, water quality, business efficiency, statutory compliance, and stakeholder confidence. The purpose of the meeting was for the CRM to obtain feedback in regard to the parameters and the severity scales, i.e. had he used the right parameters, and were the severity scales accurate?

The discussion in Dialogue 4 revealed the complex range of inter-related standards to which the company was required to perform, particularly with respect to the delivery and treatment of reticulated water. So, for instance, for water supply there were performance standards relating to factors of immediate health significance (P1 Determinands), factors of

long-term health significance (P2 Determinands), aesthetic factors (taste and odour), the quality of management systems and personnel (MoH Grading), water flow and pressure, and drought security. Evaluating the relative significance of breaches of these standards involved consideration of the importance of the standard (e.g. P1 vs P2 Determinands), the magnitude and duration of the breach, the number of people affected, control actions that the company might take (e.g. imposing water restrictions, issuing Boil Water Notices), and subsequent outcomes (e.g. people getting sick, media attention, prosecution of company personnel, restructuring of the company).

As it turned out, the CRM had identified, without too much difficulty, the key performance indicators applicable under the various standards and had even converted some of them into innovative risk assessment criteria. But he lacked the detailed operational knowledge necessary to sort the plethora of inter-related parameters into a coherent framework. That knowledge lay with the Operations management and staff who were intimately familiar with those parameters through their roles. Thus, the debate in the meeting focussed almost exclusively on the latter of the above questions, initially around the relative significance of different water supply service failures, and then on the relative significance of water supply versus wastewater service failures.

Interestingly, however, there was considerable disagreement amongst the meeting participants over how the CRM had arranged the various parameters in the risk framework, particularly with respect to the severity rankings. Indeed, although the participants made a genuine effort to arrive at an objective prioritisation of outcomes, they failed to reveal a unified view of the relative significance of different performance failures. It was clear from the discussion that each had a personal set of criteria for evaluating the significance of outcomes which was influenced not only by their specific roles and responsibilities in the company, but also by their own subjective expectations of who might be held accountable in the event that things went wrong.

For instance, a water supply grading issue was perceived to be more important to the Water Treatment Manager than the Water Networks Manager because as the former commented, “We can’t get a ‘B’ in Treatment, we can go from ‘A’ to ‘C’ to ‘D’” (Dialogue 4, p 263). Thus any potential downgrading would have a bigger impact on the Treatment Manager, in terms of accountability within the company, than on the Networks Manager. Similarly, the Wastewater Treatment Plant Manager perceived complaints about odour and midges as very significant because he felt that he was under pressure from senior management to keep those problems under control (midges were a recurrent problem at the Mangere plant because the large areas of still water provided favourable breeding conditions;

which was an issue that the \$500 million upgrade of the plant in the early 2000s had been designed to address). Ironically, while the managers from the water side of the business clearly felt that midges and odours were relatively insignificant compared with the potential effects of failures in the water supply network, they too were equally concerned with complaints about even minor water quality issues. As one manager commented, “For the last taste and odour event we had, there were 123 complaints from consumers and Metrowater were screaming. They wanted me on my knees for 123 complaints.” (Dialogue 4, p 263)

Dialogue 4 thus revealed the diversity of perceptions, across different communities of practice within the organisation, about “what mattered” with respect to risk outcomes. On reflection, the degree of variance in the managers’ revealed perceptions of the significance of risk outcomes is not surprising. Individual actors will naturally tend to prioritise outcomes that affect their area of the business and their roles, particularly if the outcomes affect areas of business performance for which they are directly responsible. It is perhaps fundamentally true of organisational life (and indeed of life in general) that no one likes to be responsible for failure, and no one likes to be told that what they do is less important than what others do. Hence, the Wastewater Treatment Plant manager was understandably indignant when he was explicitly told by the CRM during the meeting that the consequences of performance failures in his area of the business could never be as bad as the potential consequences of failures in the water supply network.

### **Engaging with stakeholders: General Managers**

The CRM initially assumed that Watercare’s Statement of Corporate Intent (SCI) was the document which contained the most definitive specification of the company’s corporate objectives. As described in Appendix I, Watercare’s primary governance instruments were in fact the Local Government Act (1974, 2002) and the Company Constitution (2004), which specified the company’s governance regime and principal objectives. The requirement to produce an SCI under the provisions of the LGA 2002 was somewhat ambiguous because Watercare was designated as a council organisation rather than a council-controlled organisation, but the company had received a legal opinion in 2003 to the effect that the company would still be required to prepare an SCI when the LGA 2002 came into effect. In any case, the Company Constitution, amended in 2004, explicitly required that the company prepare an SCI every year and within that document to specify, among other things: *the objectives of the Company* and *the performance targets and other measures by which the performance of the Company may be judged in relation to its objectives*. This requirement



meant that the SCI was intended to be the formal mechanism through which Watercare's Board and Shareholders (either in agreement with the Board or by resolution) could specify how the company was to give effect to and be judged against the principal objectives set forth in the Local Government Act and the Company Constitution. On this basis it could reasonably be argued that the SCI was formally the principal repository of enterprise objectives and key performance indicators for the company, and was therefore the logical starting point from which to derive risk assessment criteria for the CRM's revised frameworks. This view was reinforced by the company's own depiction of the corporate "management system", published in the Asset Management Plan, which positioned the SCI along with the customer contracts at the top of the system (Watercare Services Ltd. 2006, Fig. 21, p 25). In addition, the CRM was aware that the existing (2003) risk framework had been derived from the content of the SCI.

By October 2007 the CRM had produced a draft revision of Watercare's risk framework based on his deconstruction of the SCI objectives. At that time he proposed a Business Operations risk framework with five assessment categories, representing potential risk impacts on water delivery, water quality, business efficiency, statutory compliance, and stakeholder confidence. The CRM subsequently met with four of the company's General Managers to explain the content and structure of his proposed framework (Dialogue 7 is a record of that conversation).

Much to the CRM's surprise, the General Managers quickly expressed doubt about the basis of the CRM's proposed framework. The particular issue was that the CRM's framework represented Watercare's water service objectives under two categories, "delivering water" and "treating water", but did not include corresponding categories for Watercare's wastewater service objectives. Rather, the framework represented the company's wastewater service objectives under the category of "maintaining statutory compliance", i.e. a failure in the wastewater system would ultimately result in non-compliance against regulatory standards (resource consents). One manager pointed out that the same could be said about a failure on the water supply side of the business, i.e. such a failure could result in non-compliance against the Ministry of Health Drinking Water Standards. It seemed the General Managers were not comfortable with such an inconsistency, and felt that the two sides of the business should be represented consistently within the risk framework.

It subsequently emerged that Watercare's General Managers did not see the SCI as the definitive statement of the company's core objectives. One manager commented that over time the SCI had become something of a "dumping ground" (Dialogue 7, p 292) for various stakeholder demands such that it was now an unwieldy document and process to manage.

The large number of “tactical and prescriptive requirements” (Dialogue 7, p 292) in the SCI meant that, in the eyes of Watercare’s management, it was no longer a strategically focused governance mechanism. Another manager even questioned the SCI as a primary driver of the strategic plan. While the SCI “couldn’t be ignored”, it was not considered to be a “fundamental driver” of the business (Dialogue 7, p 292). Ironically, one of the general managers present in the meeting had only recently joined the company, and he commented that he too had thought the SCI “would be one of the primary drivers of everything we did” (Dialogue 7, p 293), and was therefore as surprised as the CRM to hear the other managers’ points of view.

This revelation was both a surprise and a setback for the CRM because he was forced to reconsider the basis for his risk framework. He noted a few days later (see Dialogue 8), the “interesting fact” that the SCI and the Annual Report were not consistent, each document containing different sets of objectives and performance indicators. Since the SCI and the Annual Report were the two primary documents which detailed Watercare’s corporate objectives, and given their respective positions at the beginning and end of the annual business cycle, it was to be expected that they should be reconcilable. Closer examination revealed that, in addition to having different contents, neither document clearly differentiated between statements of intent, strategic objectives, key performance indicators, or specific actions. Of the suite of performance indicators used in the 2006 Annual Report, for instance, only about half could be considered true performance indicators, the rest were simply statements of actions that the company had undertaken (i.e. they recorded what the company had done, but not the resulting changes in enterprise performance). Similarly, a number of the “objectives” in the SCI were not really objectives but were rather statements of performance standards that the company should achieve (e.g. achieving a funds flow from operations to interest cover of 3.50 times or better) or actions the company should implement (e.g. to commission a targeted cost efficiency review).

The CRM suggested in Dialogue 8 that despite their formal significance neither the SCI nor the Annual Report seemed to be fulfilling the governance functions for which they were intended. The SCI appeared to have become little more than a customer and shareholder “wishlist”, and the content of the Annual Report appeared to be driven more by “a desire to keep the report leading edge so that it continues to win awards”<sup>31</sup> (Dialogue 8, p 298), while

---

<sup>31</sup> This is a reference to the fact that Watercare’s Annual Reports consistently received awards for best practice in Sustainability Reporting, both in New Zealand and internationally.

the discord between two documents indicated that they were no longer transparently connected as a governance process. The CRM concluded that the ambiguity in the SCI, the lack of consistency between the SCI and the Annual Report, and the lack of enthusiasm from Watercare's general managers to use the SCI as the primary statement of the corporate objectives meant that it would not serve as a robust basis from which to establish the company's risk frameworks.

It is not known whether Watercare's general managers had ever explicitly acknowledged the problematic status of the SCI as a group prior to the meeting in Dialogue 7, and I did not probe further into how the situation had evolved historically because my focus at the time was to follow what the CRM was doing. While he felt the situation was cause for concern from a governance perspective, there was little he could do except to move on and try to find a solution. Indeed, the general managers also recognised that the CRM's interest in this area was a valuable opportunity to address the problem, with the expectation that he could undertake an internal review of Watercare's governing documents and recommend options for clarifying and rationalising the various versions of the enterprise objectives.

In this regard, the general managers pointed to the Watercare-specific objectives in the Local Government Act as the most fundamental and permanent statement of the company's purpose. The SCI was seen both as a translation, "even if not a very good one" (Dialogue 7, p 293), of that purpose into some practical measures of performance, and as a statement of how Watercare would contribute to the well-being<sup>32</sup> of the Auckland region. Beyond that, the managers also identified a range of forces which were important with respect to the strategic direction of the organisation, including the company's contracts with its customers, the fundamental drivers of Growth, Regulation, and Levels of Service<sup>33</sup>, the regional Three Waters Vision (2005d) and Strategic Plan (2008b)<sup>34</sup>, and the company's internal Strategic Plan and the "directors' views about where they want to take the enterprise" (Dialogue 7, p 293).

### **Explanation of the GMs' views**

Dialogue 7 leaves a number of first impressions, including that Watercare's general managers

---

<sup>32</sup> i.e. social, economic, environmental, and cultural well-being, as defined in the LGA 2002.

<sup>33</sup> Watercare's capital investment in water and wastewater infrastructure was driven by three primary forces: regional growth in water demand/wastewater volumes, changing regulatory requirements, and improvements to levels of service (usually at the request of the Local Network Operators).

<sup>34</sup> These documents were produced out of a region-wide collaborative planning programme, led by Watercare, for the integrated management of water supply, wastewater, and stormwater (the Three Waters project).

openly discounted the significance of a key governance instrument, that they lacked a unified view of the corporate objectives, or even as an illustration of the multiplicity of ways in which an organisation and its environment can be interpreted and understood by various actors, and how what counts as “fundamental” very much depends on one’s perspective and role in that system. As one general manager commented, there are “a whole lot of different ways of slicing and dicing the business, and many different answers” (Dialogue 7, p 293). However, such first impressions must be tempered in consideration of the context in which Dialogue 7 took place.

Watercare was a public sector service organisation, wholly owned by the local authorities of the Auckland region. The principal objectives of the company had literally been “set in stone” by the Local Government Act and had remained unchanged since the company was formed in the early 1990s. The company’s principal objective, as stipulated in that Act was as follows: “Watercare Services Limited must manage its business efficiently with a view to maintaining prices for water and wastewater services at the minimum levels consistent with the effective conduct of that business and the maintenance of the long-term integrity of its assets.” The Act essentially stated that Watercare existed to provide bulk water and wastewater services to the Auckland region, and that the company’s principal objective in fulfilling that purpose was to balance the long-term costs of its services against the long-term risks to the sustainability of those services. As a public sector organisation governed by an Act of Parliament, Watercare and its officers were not at liberty to choose to alter the company’s purpose or core objectives, or the range of business the company was to pursue. Consequently, corporate strategy at Watercare was limited to decisions about the policies and plans for achieving its goals, and about the kind of organisation that the company aspired to be. In other words, Watercare’s strategy was limited solely to the question of how the company should fulfil its designated purpose (see the later discussion over what “strategy” meant in Watercare’s context).

In this regard, the Company Constitution specified additional objectives. That is, the company’s services should be “economically viable, environmentally sound, socially responsible and responsive to customer needs.” However, as I noted in Dialogue 8, the statements in the Company Constitution were extremely general. The Constitution was not specific as to what “environmentally sound” or “socially responsible” meant in practice. But nor would it have been appropriate for the Constitution to be specific in that regard. This was, rather, the formally designated purpose of the SCI, and, indeed, this is what the SCI did. Even if the document could be criticised as ambiguous, unwieldy, or lacking in strategic focus, it constituted, in effect, a statement of Watercare’s corporate social responsibility

(CSR) policy. The SCI described what sort of organisation the company should aspire to be, and how it should operate vis-à-vis the delivery of its services. That is, Watercare was to be an organisation that maintained a dialogue with stakeholders, customers, and the community, co-ordinated planning activities with its customers, sought out opportunities to reduce its environmental footprint, and to improve the efficiency of its operations, aspired to best practice in asset management and employment relations, provided its employees with safe working conditions, adhered to fiscally prudent policies, and otherwise complied with all statutory and regulatory obligations applicable to its business.

Dialogue 7 revealed that Watercare's general managers did not perceive the SCI to be a strategically relevant document. This is not to say that the SCI was unimportant. On the contrary, Watercare was well respected as a responsible corporate citizen, indeed a corporate leader, precisely because the company operated in accordance the SCI. Rather, it was the case that the SCI had little impact with respect to the strategic development of the water and wastewater infrastructure. The reason for this was that Watercare's strategic decision making was constrained by a physical and institutional context that was not only rigid, but static<sup>35</sup>. The physical environment within which the company operated was defined largely by a geographically extensive water and wastewater infrastructure, which represented over 100 years of prior capital investment and maintenance. Decisions about how to deliver water and wastewater services were significantly constrained by the historical development trajectory of that infrastructure system and high sunk costs of prior investment decisions; both typical features of Large Technical Systems (Hughes 1987; Joerges 1988; Katko et al. 2006; Monstadt and Naumann 2005). In this context, the strategic development of the system was the product of a supply – demand equation, where “demand” encompassed anticipated future production, changes to levels of service and regulatory conditions, and anticipated greenfield developments in the city, and “supply” encompassed the capacity and reliability of the existing infrastructure (which changes over time as assets deteriorate), the evolving state of technology, and the resources available to construct new infrastructure, and renew or replace existing infrastructure. The two sides of the equation were represented by the three forces referred to in Dialogue 7: Growth, Levels of Service, and Renewal. Also mentioned in Dialogue 7 was the Three Waters Vision, which was strategically important as a joint planning exercise with the Local Network Operators to identify efficiencies from the

---

<sup>35</sup> Watercare's operating context and its influence on the technological development of the water and wastewater infrastructure are elaborated further in Appendices I and II.

integrated planning of water, wastewater, and stormwater works across the region. The customer contracts with the LNOs, also mentioned, were strategically important with respect to defining certain levels of service. What the CRM encountered in Dialogue 7 was the general managers' finely tuned understanding of Watercare's strategic context, and how it related to their respective areas of responsibility.

## **The revised (2008) Operations Risk Framework**

Following his engagement with Watercare's General Managers (described in the previous section), the CRM recognised that he would not be able to derive his revised Operations Risk Framework from the company's Statement of Corporate Intent. Rather than the SCI, the CRM's took the Local Government Act (1974, 2002) and the Company Constitution (2004) as the primary documents of reference for defining the corporate objectives. The general statements of purpose contained within those documents were explained in Appendix I. The framework that the CRM subsequently proposed, in early 2008, to replace the existing (2003) framework is summarised in Table VI.1.

Table VI.1 shows how the CRM deconstructed Watercare's principal objectives into ten fundamental performance objectives for the company (left hand column). Rather than attempt to arbitrarily define the criteria by which the successful achievement of the corporate objectives should be judged, the CRM explicitly looked for and appropriated existing performance standards and criteria (middle column in Table VI.1). As described earlier, he had already identified that there were a number of external standards to which the company was required to respond (statutory obligations and the customer contracts), and that the company's performance was also regulated by a number of internal standards (the operational budget, the Asset Management Plan, and the Funding Plan). To these the CRM added two additional standards, the level of customer (i.e. LNO) satisfaction with Watercare's performance (defined in different terms than the 2003 framework), and the level of "unacceptable" risk carried by the company on a long-term basis for a given level of expenditure; the latter being an indicator of performance against the company's core strategic objective to balance the long-term cost of its services against the long-term risk to the provision of those services. The CRM divided the performance indicators into two categories, representing measures of short-term (day-to-day) performance, and measures of long-term (year-on-year) performance (right hand column in Table VI.1). Those general measures were then further developed into functionally specific evaluation categories and criteria for the Operations framework.

**Table VI.1. Measuring performance against corporate objectives**

<b>Corporate Objectives</b> <i>Local Government Act 1974</i> <i>Company Constitution 2004</i>	<b>Performance Measures</b> <i>(i.e. performance against the corporate objectives may be measured by...)</i>	<b>Performance objectives for the risk framework</b>
Manage the business efficiently	The degree to which the organisation achieves the budget.	Short term (day to day) measures of business performance:  1. Achievement of the budget  2. Compliance with customer contracts  3. Compliance with statutory obligations  4. Customer satisfaction  Long term (year on year) measures of business performance:  5. Achievement of the Asset Management Plan  6. Achievement of the Funding Plan  7. Level of “unacceptable” risk carried by the organisation.
Provide water and wastewater services	The degree of compliance with customer contracts and statutory obligations (particularly DWSNZ and MoH Grading requirements for water supply)	
Operate as a successful business	The degree to which the organisation achieves the Asset Management Plan and the Funding Plan.	
Provide services that are economically viable	The degree to which the organisation achieves the budget and Funding Plan.	
Provide services that are environmentally sound	The degree of compliance with statutory obligations (esp. resource consents for the treatment plants and wastewater network)	
Provide services that are socially responsible	The degree of compliance with statutory obligations (especially obligations with respect to employment, human relations, health and safety, and reporting)	
Provide services that are responsive to customer needs	The degree of customer (LNO) satisfaction with WSL performance.	
Maintain minimum prices for services	The degree to which the organisation carries an “unacceptable” level of risk for a given level of expenditure.	
Maintain the effective conduct of the business		
Maintain the long-term integrity of its assets		

Sources for Table VI.1: (Clement 2008a, 2008f)

### **Justifying the revised (2008) Operations Risk Framework**

The framework described in Table VI.1 was to my knowledge at least the fourth variation of an enterprise-level system for the evaluation and representation of Watercare’s performance.

Three earlier variations included:

- *The existing Statement of Corporate Intent and Annual Report, and existing (2003) risk framework:* These three documents constituted a system for the representation of Watercare’s performance which had been evolving since the company’s formation in

the early 1990s.

- *The Planning Support Framework and revised risk framework (2006)*: In 2005-06 I was involved in a project to develop an integrated enterprise performance evaluation framework and revised risk framework for the company (described in Appendix II). That framework was intended to achieve many of the same outcomes that the CRM later pursued.
- *The CRM's revised risk framework (2007)*: The aborted (2007) framework that the CRM initially derived from the Statement of Corporate Intent was ultimately not that different from the subsequent (2008) framework derived directly from the LGA and the Company Constitution. Nevertheless, the starting points and the deconstructive processes for each framework were quite different.

Although there were, of course, similarities in structure and content between the above frameworks, each constituted a different representation of the Watercare universe. Since the deconstructive process for each of the four frameworks was different, together they constituted a plurality of possible, plausible interpretations of Watercare's corporate objectives. Such a multiplicity of possibilities for representing and measuring organisational performance raises an important question: how to justify the choice of any particular system as "better" than the others. This was a significant question for the CRM, particularly because there was some scepticism from Watercare's managers over the need for changes of the scope and magnitude that the CRM intended to make to the company's risk framework. Since the existing framework had been in use for many years, and had originally been developed under the supervision of one of the existing General Managers, the CRM felt that he needed to have strong case for change.

The case that the CRM put forward was that his 2008 framework offered a more objective way of looking at the business, and would thus produce better (i.e. more objective) risk information than the existing 2003 risk framework. He argued that the new framework was more objective because it more effectively married together the strategic and operational levels of the organisation (see Dialogue 9). From a top-down perspective, the new framework established an explicit link between Watercare's corporate objectives and the various real standards which defined both the short and long-term performance of the organisation. From a bottom-up perspective, this link also served to resolve the terminology problem described in Chapter 3. By referencing the actual performance standards which drove the organisation, the CRM was able, for the most part, to quantitatively express the risk



evaluation criteria in specific operational terms. In the CRM's own words, he was able to express the framework in "language that the guys at the coal face use on a day to day basis" (see Dialogue 9, p 305).

The CRM's expectation was that the new framework would allow for more objective risk assessments than the existing framework. First, by explicitly defining risk in terms of actual performance standards, the CRM's framework eliminated the need for staff to subjectively assess the follow-on impacts of poor performance, thus improving both the absolute and ethical objectivity of the framework. Second, because the language of the framework was more explicitly and transparently linked to the actual work that the company performed, operational knowledge could now be represented using terms and measures that were concrete to the functions of the business. This latter point also offered the benefit that the criteria and measures within the framework could be quantified or otherwise supported by real data from Watercare's existing business systems. The CRM believed that this would be a major benefit in support of his framework, particularly the possibility of quantifying risk using existing modelling capabilities (e.g. the RCM models). The CRM's expectation was that these improvements would constrain the subjectivity inherent to the translation of information from the bottom to the top of the organisation, thus resulting in the production of more objective representations of risk, and making communication between "the coal-face and the Board room" more efficient and effective (see Figure 4.1, Chapter 4, p 101)

## **Redesigning the Project Risk Register**

After redefining Watercare's risk assessment frameworks the Corporate Risk Manager undertook to redesign the company's risk registers. This was necessary to ensure that the registers were consistent with the new frameworks. The following sections briefly summarise the design of the new Project Risk Register and the surprising difficulties that were encountered when it was put to use in a project risk identification workshop. The implications of this outcome are discussed in Chapter 4 of the thesis.

### **The new Project Risk Register**

As was described in Chapter 3, the CRM felt that risk descriptions in the company's existing risk registers were often ambiguous, incomplete, overlapping, or simply too brief. Improving the specification of risks in those registers was an important component of the CRM's plan to

generate good quality “risk data”. The CRM therefore undertook to re-design the company’s corporate and project risk registers to reflect the revised frameworks, and to provide customised tools for capturing statements of risk in the appropriate formats.

The CRM’s approach to redefining the company’s risk frameworks was described in Chapter 4 and at the beginning of this appendix. Briefly, his objective was to specify the risk assessment parameters in terms familiar to the project managers while also maintaining the link with the corporate objectives. To achieve this the CRM met with general and project managers from Watercare’s Asset Management group in order to understand (a) the processes and work flows within the Projects team, (b) the relationships between Projects, Planning, and Operations, and (c) how the broader Watercare enterprise actually “felt” the impacts of “cock-ups” by the Projects team. The objective of the CRM’s investigations was to make explicit the context within which the project managers were working, and what the impacts of project-level risks were on the Watercare enterprise as a whole. Those discussions (see Dialogues 5 and 6) and the issues to which they gave rise were summarised earlier.

The project risk register which the CRM sought to replace had been in use within the Watercare for several years and was structured around the company’s 2003, one-size-fits-all framework. The categories and criteria for evaluating the consequences of project risks were thus the same as for evaluating all other types of risk in the company, with the addition of a separate category to record the magnitude of impacts on the project schedule (see Table VI.2). Through his investigations the CRM identified that there were five primary objectives applicable to projects: (1) to satisfy a functional output objective – i.e. to deliver a certain functionality or business capability; (2) to satisfy a schedule objective – i.e. to deliver that capability within a specified time frame; (3) to satisfy a budget objective – i.e. to deliver that capability within a specified budget; (4) to comply with all statutory requirements; and (5) to impose minimal interference on stakeholders (both internal and external to the Watercare organisation). These objectives subsequently became the assessment categories in the CRM’s new project risk register (see Table VI.2).

In order to calculate the magnitude of risk consequences relative to the project in question, the user of the register was required to specify the particular project objectives: (1) the criticality of the functional output, (2) the budget, and (3) the targeted and critical completion dates. Formally, the “criticality” of the project could be measured by the magnitude of the risk the project was designed to mitigate, but to keep things simple the register required the user only to rank the importance of the project using a basic qualitative scale. The critical completion date was the CRM’s method of accounting for the fact that even substantial schedule delays could often be accommodated without significant impacts

on enterprise performance. The register required the user to specify a critical completion date as the estimated absolute latest date at which the project could be completed before the delay would adversely affect enterprise performance. When assessing risks, the user was required to qualify or quantify the expected consequences in terms of the degree to which the functional output would be affected, the additional financial expenditure required, and the anticipated delay to the project programme. The register then calculated two risk scores, the Project-level score and the Corporate-level score, as shown in Table VI.2; the former representing the significance of the risk for the project manager, and the latter representing the significance of the risk to the Watercare enterprise. The CRM felt that the new project risk framework and register addressed two important needs: (i) that of providing project managers with a tool which would enable them to assess risk in terms relevant to their work context, while (ii) also enabling project level risks to be objectively represented in the corporate context.

*Table VI.2. Comparison of the evaluation categories in the existing 2003 and revised 2007 project risk registers*

<b>Evaluation categories in the existing 2003 project risk register</b>	<b>Evaluation categories in the revised 2007 project risk register (also showing calculation of Project vs Corporate Risk Profiles)</b>
Programme delay	Cost overruns, calculated as follows: <ul style="list-style-type: none"> <li>▪ <u>Project Profile</u>: Additional expenditure as percentage of original project budget.</li> <li>▪ <u>Corporate Profile</u>: Additional expenditure.</li> </ul>
Reputation	Schedule overruns, calculated as follows: <ul style="list-style-type: none"> <li>▪ <u>Project Profile</u>: Delay as a percentage of the original project programme (i.e. targeted completion date).</li> <li>▪ <u>Corporate Profile</u>: Delay as percentage of time available until the critical completion date is reached</li> </ul>
Finance	Functional output, calculated as follows: <ul style="list-style-type: none"> <li>▪ <u>Project Profile</u>: Magnitude of impact on functional output relative to original specification.</li> <li>▪ <u>Corporate Profile</u>: Magnitude of impact on functional output relative to original specification, multiplied by the criticality factor.</li> </ul>
Environment and Public Health	Statutory compliance (Project Profile same as Corporate Profile)
Health and Safety	Impact on stakeholders (Project Profile same as Corporate Profile)
Asset Management	

Sources for Table VI.2: (Clement 2007b; Watercare Services Ltd. 2003c)

## **Surprising difficulties using the new register**

Around 75% of the treated water supplied by Watercare to the Auckland Region was derived from Watercare's southern water sources: four dams in the Hunua Ranges which delivered raw water to the Ardmore Treatment Plant, and the bulk treated water pipeline from the Waikato River Treatment Plant. All of this water passed through a large reservoir complex at Redoubt Road, approximately 20 km southeast of the Auckland CBD. From the Redoubt Road Reservoirs, the treated water was transmitted to the city through three large water mains: the Hunua No. 1, No. 2, and No. 3 mains; of these the Hunua No. 3 (HN3) main was the largest, delivering approximately 50% of the city supply. During 2007 Watercare undertook detailed planning for the construction of another large treated water main from the Redoubt Road complex into the city, to be called the Hunua No. 4 (HN4). This main was required to meet long-term future growth in demand, but, more immediately, was required for operability reasons so that the HN3 main could be taken out of service for maintenance. The estimated capital cost of the HN4 main was NZ\$195 million.

By December 2007 the HN4 project was nearing the end of the detailed planning phase. At this point, the route for the pipeline, and the corridor within which the pipe was to be laid along the length of that route, had been finalised. Although important details such as the size of the pipe had already been established, other key decisions, such as about the type of lining material, the precise placement of the pipe within the design corridor, and programme scheduling, were to be completed during the detailed design phase. In preparation for the delivery phase of the project (i.e. detailed design and construction), a workshop was held to identify key risks that would need to be taken into account during that phase. The objective of the workshop was to identify risks that the design consultant would need to address in the design of the pipeline, and for which the construction contractor(s) would need to provide management plans. The workshop was attended by Watercare's Risk Specialist, two project managers, and myself. In the normal course of events the workshop would have been facilitated by the Corporate Risk Manager, but he could not attend due to a scheduling conflict. The workshop was additionally significant because it was the first use of the CRM's new Project risk framework and register. I performed the role of scribe, using the new risk register as a template to record the "risk data" which emerged from the workshop, but I did not participate in the discussion (other than to clarify certain points from time to time). The workshop was not audio recorded; the account presented here was reconstructed from my notes made during and immediately after the workshop.

The Risk Specialist kicked off the workshop with a list of potential risks/issues that he

had previously brainstormed, and the group proceeded to discuss each issue in order. The discussion generally took the form of an open forum with little in the way of a formalised approach to defining a statement of each risk. In most cases, only the generic nature of the potential consequences of a risk (e.g. schedule delay, or operational impact) were identified, sometimes with a generic comment as to the significance (e.g. “this is a show stopper”), but the specific likelihood and consequences were not otherwise made explicit in terms of the Project Risk Framework. In contrast, the discussion of the actions and controls needed to address each risk/issue were usually quite detailed. Indeed, the discussion tended to focus very quickly on what issues needed to be addressed, various options for resolving the issues (e.g. design alternatives), and future work to be done (e.g. additional investigation, modelling etc.). These included, for example, specific technical options or approaches, identification of specific stakeholders to be consulted, specific problems to be anticipated, brainstorming of innovative solutions to be investigated. In this sense, the definition of the identified risks was considerably lopsided, with little attention paid to clearly defining the specific risk event, its causes, or its consequences, but significant detail emerging in the discussion of the risk controls.

In quite a number of cases, the discussion centred around an uncertainty in the project plan rather than a specific risk. For example, at that time the options for a particular river crossing were not finalised. There were several potential options on the table, and the choice of option was contingent on the outcomes of further design work and consultations with external parties. In this and other similar cases the discussion centred on what needed to be done to resolve the uncertainty and finalise the design, but the uncertainties were not discussed in specific risk terms (i.e. cause, event, consequences). The discussion also addressed opportunities for operational improvements or co-ordination with other projects as often as potential downside risks.

I found the workshop to be particularly interesting because it turned out that the newly designed project risk register appeared to be not well suited to capturing the outcomes of the workshop. The register required the user to specify the overarching project objectives: the project deadline, the estimated budget, and the nature of the deliverable or expected functional output. These then served as a baseline for quantifying the potential consequences of risk events, i.e. the degree of deviation from the schedule, budget, or deliverable specifications. It was expected that this systematic way of defining the project management context would facilitate the clear and unambiguous description of risks. However, as scribe, I found it difficult to format the emerging information to the structure of the risk register template. Certainly, this was partly due to the relatively fluid and unstructured nature of the

discussions; had the CRM been in attendance, he would likely have facilitated a more structured approach, which would have made the data entry task somewhat easier. However, even after collating and rationalising the information from the workshop, I found that the risk register could only be incompletely populated. Since the new Project Risk Framework had been carefully designed to facilitate the definition of risks in the project management context, and since I was reasonably confident that the problem did not lie with my ability to interpret and record the outputs of the workshop conversation, I was faced with a perplexing question: why had such a carefully designed tool proven so difficult to use in the very context for which it was intended?

## **The availability and quality of data for risk modelling**

The CRM promoted quantitative risk modelling because it offered, at least in the ideal, the means to objectively evaluate the aggregate risk exposure across groups of related risks. Improved insight into the company's aggregate risk exposure would lead to better prediction of asset replacement and renewal requirements, provide a more analytical basis for judgements of risk tolerance and prioritisation, and, hence, provide more objective support for capital expenditure decisions and the preparation of the Asset Management Plan (see Dialogues 1, 2, and 3). Quantitative modelling also constituted a potentially important selling point in support of the CRM's revised risk frameworks (see Dialogue 9).

Detailed reliability models had already been developed for the company's water and wastewater treatment plants. For the CRM, the next logical step was to extend the modelling programme to the rest of Watercare's physical infrastructure (i.e. to the water distribution and wastewater collection networks) and ultimately to employ quantitative probabilistic modelling to aggregate all of the company's risks into a single Enterprise Risk Profile. The CRM expressed this vision in Dialogue 2.

Briefly, the CRM envisioned that asset condition and reliability information could be correlated with asset age, and, eventually, with "good" data, it would be possible to establish a "more analytical" basis for assessments of the likelihood of asset failures. Combined with assessments of the consequences of failure, this would allow the risk of asset failures to be quantified with "greater certainty" than at present. With such models systematically defined, the CRM envisaged that it would subsequently be possible to obtain robust, quantitative estimates of Watercare's aggregate exposure across groups of risks. So, for instance, it would become possible to answer such questions as: for example, for the water supply network,

what would be the likelihood of a failure event resulting in loss of supply to more than 50,000 people for more than 8 hours in the next 12 months? Or, across Watercare's programme of capital works, what would be the likelihood of a financial loss greater than \$10 million in the next 12 months? Or, for the enterprise as a whole, what would be the expected value for Capex and Opex variance over the next 12 months? Such quantities could then be compared with pre-defined tolerance criteria to objectively establish whether the risk was unacceptable, tolerable, or acceptable.

However, since the data requirements for reliability modelling are significant, the CRM asked me, in early 2008, to review the availability and quality of Watercare's existing asset condition information for the company's water and wastewater network assets. The task was essentially to conduct an assessment of the existing data quality and comprehensiveness as a baseline for prioritising a data quality improvement programme. I performed the task by trawling through Watercare's networked computer drives and various electronic databases, reviewing earlier reports on the same subject, and talking to a number of people from the Operations group. As a result of that investigation I came to the understanding that Watercare's asset condition information did not exist as a cohesive, consistent, or even singular data set. Rather, there was no central repository of asset condition information and what data did exist was widely variable in terms of comprehensiveness and quality, ranging from reasonably complete, consistent, and reliable data for some assets, through to virtually non-existent, incomplete, or highly subjective data in others.

Asset maintenance and repair histories could be established for any specific asset from work order logs recorded in Mozaic (the primary business financial information system), and failure histories were typically recorded in independent Excel® spreadsheets or Access® databases, but there was no central database which recorded all of this information. Rather condition information was distributed throughout the company in a number of databases, in electronic files in a myriad of locations on the company network, and in various physical files located in filing cabinets throughout the company offices, the company archives, or simply on the desks of individual employees. Information was also available in the form of the personal knowledge and experience of company personnel. The completeness and quality of information available for different types of assets was also highly variable. For some assets there were complete maintenance and repair histories, failure records, and regular condition inspection reports, while for others the most recent records were as-built drawings, sometimes decades old.

In theory, asset condition information is a critical component of both maintenance and capital planning calculations, and, as such, the relative absence and disorganisation of this

information for a significant portion of Watercare's assets was initially disconcerting. However, while there was indeed room for improvement<sup>36</sup>, the range of data sources and qualities should not have been unexpected for several reasons. First, condition inspection and assessment can only be cost effectively carried out on accessible assets. Thus, for instance, Watercare had virtually complete CCTV coverage of the company's sewer network because sewers are accessible during periods of low flow, but there was very little condition information available for watermains because mains flow under pressure and are relatively inaccessible (they must be shut off and drained in order for an inspection to be carried out). Available condition information for assets in the water supply distribution network was limited to above ground assets, or to isolated sections of underground pipe which had been opportunistically inspected as a result of a failure (i.e. at the time the pipe was repaired). Second, forward maintenance planning is simply not required for some assets because the consequences of their failure are negligible. Such assets can generally be run to failure and knowledge of their condition is therefore not required. And third, the expected productive lifetime of some assets is very long, on the order of 50 – 100 years. For such assets it is reasonable to expect that the company would not need to worry about the condition of the assets for many years after their construction, unless there were indications to the contrary (e.g. unexpected failures).

## **Facilitating the identification of strategic risks**

Workshop facilitation was not identified as an explicit element of the Corporate Risk Manager's vision and plan for developing Watercare's risk management capabilities, except in the context of training. That is, the CRM anticipated a need to educate management and staff on various risk management practices, but did not otherwise identify facilitation as a significant component of the role. Nevertheless, facilitating risk identification and assessment workshops did become an important aspect of his work. Typically this involved guiding other actors through an explicit and systematic process of context definition, risk identification, and assessment; sometimes spread over several workshops, depending on the complexity of the system in question.

Although the CRM was himself a professional engineer, his background was in fire engineering not water supply and wastewater engineering. Consequently, in most cases, the

---

<sup>36</sup> Watercare was proactive in this regard, initiating an Asset Data Quality Improvement (ADQI) project in 2007.



CRM was not in a position to contribute specific operational knowledge to the discussion. Rather, the CRM would typically focus on leading the workshop, presenting the problem context, and moderating the discussion to keep it focussed and on track (this often being more difficult than it first seemed given the tendency of engineers to jump into the detail of problem solving).

A risk identification workshop would typically proceed as follows. In the first phase of a workshop, a considerable amount of time was often devoted to setting up the problem frame and to introducing the knowledge that was necessary to approach the problem. This was usually achieved with the aid of a Powerpoint® presentation. During this time the workshop participants were free to ask a lot of questions, mainly probing the basis of the design scenario. Since risk in Watercare's context was most often associated with the failure of physical assets in the water and wastewater networks, the risk context was usually defined in terms of some component of those networks (e.g. a pump station, a reservoir, etc.), and the relevant service level requirements. Following the problem definition phase of the workshop, attention would centre on the system in question with considerable time devoted to ensuring that the workshop participants clearly understood the layout, function, and capacity of the various assets within that system, as well as the degree of connectivity and resilience between those assets. This usually took the form of a systematic question and answer session on each asset (e.g. "What happens if we lose this asset? Can we still meet the supply conditions?"). The discussion would be supported by various large scale (A0) maps, plans, and aerial photographs of the system, and experienced technical specialists from the Operations group would be in attendance to provide specific and detailed answers to participants' questions. Identified risk events were recorded for later analysis and assessment.

Perhaps the most significant of the CRM's facilitative efforts was the work that he undertook to define a strategic risk framework for the company and to elicit knowledge of strategic risks from Watercare's managers. This is described in the following sections.

## **Identifying strategic risks**

In September 2007 the CRM asked middle and senior managers to "consider and describe emerging trends or foreseeable future conditions that could have a major impact on how WSL does business in the medium to long term (i.e., 10 to 50 years)." (Email from CRM to Managers, September 2007). These external environmental factors would be considered the organisation's strategic risks. The CRM asked the managers to categorise and record their

thoughts on a standard PESTEL template<sup>37</sup>. In December 2007 the CRM repeated the exercise, but this time focussing on internal sources of risk, that is risks “arising from the way that WSL conducts its business” (Clement 2008e). The managers were asked to categorise and record their thoughts according to the following template: Business Structure, Personnel, Process, Technology, Infrastructure. The CRM did not individually or collectively interview the managers, but rather requested that they should record their thoughts on the predefined templates over a period of approximately one week.

The CRM ask me to collate the responses from the individual managers and compile a summary report. Dialogue 10 is a record of the conversation between myself and the CRM when we met to discuss the results of that exercise. I reported that although the manager’s had identified a broad range of issues that they felt were potential strategic risks, their specification of those issues was, almost without exception, particularly vague. They tended to be little more than short sentences, or even, in many cases, simply one or two words, e.g. “terrorism” or “tightening of standards”. Virtually no indication was given as to either (a) why a particular issue was ‘on the radar’ or (b) what the impacts of that issue might be for Watercare. So, for example, was there current evidence of a heightened terrorist alert in the Auckland region? Which standards were the managers worried about and why were they concerned that those standards would become more stringent? And what would be the implications of more stringent standards? In addition, different managers had identified and categorised the same issue in different ways. For instance, similar governance issues were classified by various managers under three different categories, Political, Economic, and Legal/Regulatory.

Although initially frustrating and surprising to both myself and the CRM, the nature of the managers’ responses was, on reflection, not unexpected. First, managers are busy and it is simply easier to jot down a short prompting statement rather than a long, detailed explanation. This suggests, second, that the nature of the manager’s responses was at least partially a function of the method of data capture. The responses would likely have been more detailed and complete if the CRM had undertaken to interview each manager individually. And, third, given the scope and complexity of the issues identified, it could not be expected that the managers would possess comprehensive knowledge of every issue they thought of, in which case the vagueness of the responses could be attributed to the

---

<sup>37</sup> Political, Economic, Socio-cultural, Technological, Environmental, and Legal/Regulatory (see PEST analysis in Wikipedia, [http://en.wikipedia.org/wiki/PEST\\_analysis](http://en.wikipedia.org/wiki/PEST_analysis)).

underlying uncertainty surrounding the issues in question.

### **Defining the term “strategic risk” in Watercare’s context**

In total, Watercare’s managers had identified over 330 strategic issues (203 external, 135 internal). The CRM recognised that the managers’ responses effectively represented the results of a relatively unstructured brainstorming session, and that further work would be required to rationalise the list and extract more detailed information about the perceived risks. The CRM wished to facilitate a more structured and systematic approach to the identification and analysis of strategic risks, and thus began to think about how to define Watercare’s strategic context: “if I want to identify strategic risks then I have to ask “what are we trying to achieve at a strategic level?” and then “what could stop us achieving that?”” (Dialogue 10, p 309). The discussion in Dialogue 10 subsequently turned to the question of what the term “strategic” meant in Watercare’s context.

One general definition is that corporate strategy is “the pattern of decisions in a company that determines and reveals its objectives, purposes, or goals, produces the principal policies and plans for achieving those goals, and defines the range of business the company is to pursue, the kind of economic and human organization it is or intends to be, and the nature of the economic and noneconomic contribution it intends to make to its shareholders, employees, customers, and communities” (Andrews 1980, p 18). The CRM noted, however, that Watercare could not strategically alter the nature of its core business. The company operated as a monopoly wholesaler and was subject to a unique governance regime positioned somewhere between the private and public sectors. Although Watercare was a private limited liability company operating under the governance of an independent board of directors, it was wholly owned by the local government authorities of the Auckland region, and its existence, purpose and scope of operations were defined by legislative mandate (see Appendix I). Due to the public health significance of the company’s services, Watercare was also required to be responsive to a complex range of regulatory drivers (including the New Zealand Drinking Water Standards, resource consents, and customer contract performance requirements). Thus, unlike a fully private enterprise, Watercare could not choose to alter its purpose or core objectives, or the range of business the company was to pursue (i.e. the company could not develop new products or expand into new markets). Corporate strategy at Watercare was therefore limited to decisions about the policies and plans for achieving its

goals, and about the kind of organisation that the company aspired to be (i.e. about how the company went about its business)<sup>38</sup>.

But even in this regard the company faced significant technological and economic constraints rooted in the nature of the physical water and wastewater infrastructure. Such large technical systems (LTS) are characterised by high capital intensity, long planning and pay back periods, high sunk costs of investment, and path dependency of investment decisions with significant potential for certain strategic decisions to orient the future operation and development of the system for years to come; all of which creates significant inertia to change due to the sheer mass of technological and organisational system components (Hughes 1987; Joerges 1988; Katko et al. 2006; Monstadt and Naumann 2006). In other words, decisions about how to deliver water and wastewater services in the future would always be significantly constrained by the historical technological trajectory and sunk costs of the existing system.

The full implications of the above distinction did not become evident until the CRM and I undertook to organise and rationalise the large number of strategic issues identified by Watercare's managers in January 2008. We did this as a one-day workshop where we focussed first on establishing a framework for defining and categorising strategic risks, and subsequently on rationalising the list of strategic issues within that framework. It seemed initially self-evident that strategic risks should be defined in terms of the potential to adversely impact Watercare's *long-term* performance against its strategic objective to maintain a balance between expenditure and risk, as specified in the Local Government Act. Thus a strategic risk was one which would result in one or more of the following outcomes (notes from workshop, January 2008):

- i) A systemic and sustained upward trend in the long-term cost of the company's operations (which would translate into sustained increases in the cost of water and wastewater services, thus constituting a failure of the company to maintain prices for water and wastewater services at minimum levels);
- ii) A systemic and sustained decrease in the reliability of the company's performance against service delivery standards (constituting a failure of the company to maintain

---

<sup>38</sup> In Dialogue 10, the metaphor of climbing a tree to see the way ahead when lost in the bush was helpful in grasping this distinction: "it seems there's two levels of strategy there, one is the objective [which direction to go], and one is the things we need to take into account in order to get there [which path to take, obstacles to avoid etc.]" (Dialogue 10, p 312).

the effective conduct of the business);

- iii) A systemic and sustained decrease in the reliability of the company's physical assets (constituting a failure of the company to maintain the long-term integrity of its assets, and which would manifest in the form of (ii) above).

While operational and project risks would also adversely impact the reliability of the company's assets, as well as its financial and service delivery performance, we initially distinguished operational risks (as a conceptual category) on the basis that the impacts from operational risks would be isolated, limited in scope, and relatively temporary. That is, they would temporarily impact performance against a single performance standard, or a single asset or sub-system, and result in relatively minor unplanned expenditure. In contrast, the potential impacts of strategic risks would be widespread and sustained, constituting not just a temporary performance interruption but a persistent capability failure. Our initial distinction between operational and strategic risk was therefore one of the magnitude of the impacts.

However, we subsequently realised that the above distinction might not be sufficient to fully define the boundary between strategic and operational risks in Watercare's context. We considered the possibility of a catastrophic loss of a significant asset or sub-system (what were already referred to within the company as Strategic Assets). The collapse or failure of one of the dams impounding Watercare's raw water reservoirs was a case example. Such an event would probably result in all of the outcomes listed above, and those outcomes would be systemic in nature and persistent for a long period of time. Thus, on the basis of the magnitude of the impacts alone, such an event would have to be considered a strategic risk. But while such an event would indeed have major and lasting consequences for the company, the company also already possessed the necessary capabilities to address and control those risks (at least to the extent that they could be addressed and controlled prior to the fact). That is, they were already controlled by existing business processes; a fact which suggested they would be better classified as simply large Operational risks.

It was on the basis of the above distinction that the CRM subsequently defined a strategic risk as any event or scenario which threatened the ability of Watercare to fulfil its statutory purpose under the Local Government Act, *and* which required the company to change the way it did business in order to avoid, accommodate, or cope with the potential adverse implications of the identified scenario (Clement 2008e). This compound definition provided the necessary criteria by which to distinguish strategic risks from operational and project risks in Watercare's context, and to evaluate the magnitude of those risks. Strategic risks

were those that required the company to implement a strategic response, being one which resulted in the company developing a new capability. Operational risks were those for which the company already possessed the necessary competencies for control.

## **Rationalising and categorising strategic uncertainties**

Having reached an understanding on the definition of “strategic risk”, the CRM and I turned our attention to sorting through the large list of issues identified by Watercare’s managers. For this task we choose to abandon the frameworks originally used to capture the strategic issue data<sup>39</sup> when it became evident that there would be too much subjectivity involved in attempting to sort the strategic issues into those categories (there was considerable latitude for interpretation, and many of the identified issues could be located across multiple categories). Instead, we started with a very simple internal/external systems model representing the company, its environment, and its internal processes. The “external” category represented all those factors located in the external environment which could potentially influence the company’s performance by impacting connections between the company and its environment. The “internal” category represented those factors internal to the company which could potentially influence the company’s ability to perform its capabilities. In each case, the resulting sub-categories were produced by collating the strategic issues identified by Watercare’s managers and aggregating them together according to common underlying sources of uncertainty. Figure VI.2 shows the final collation of strategic sources of uncertainty.

We employed a crude form of grounded methodology, allowing the categories to emerge from the data itself, rather than imposing predefined categories on the data. Indeed, various commonalities were evident from very early on, as I commented in Dialogue 10:

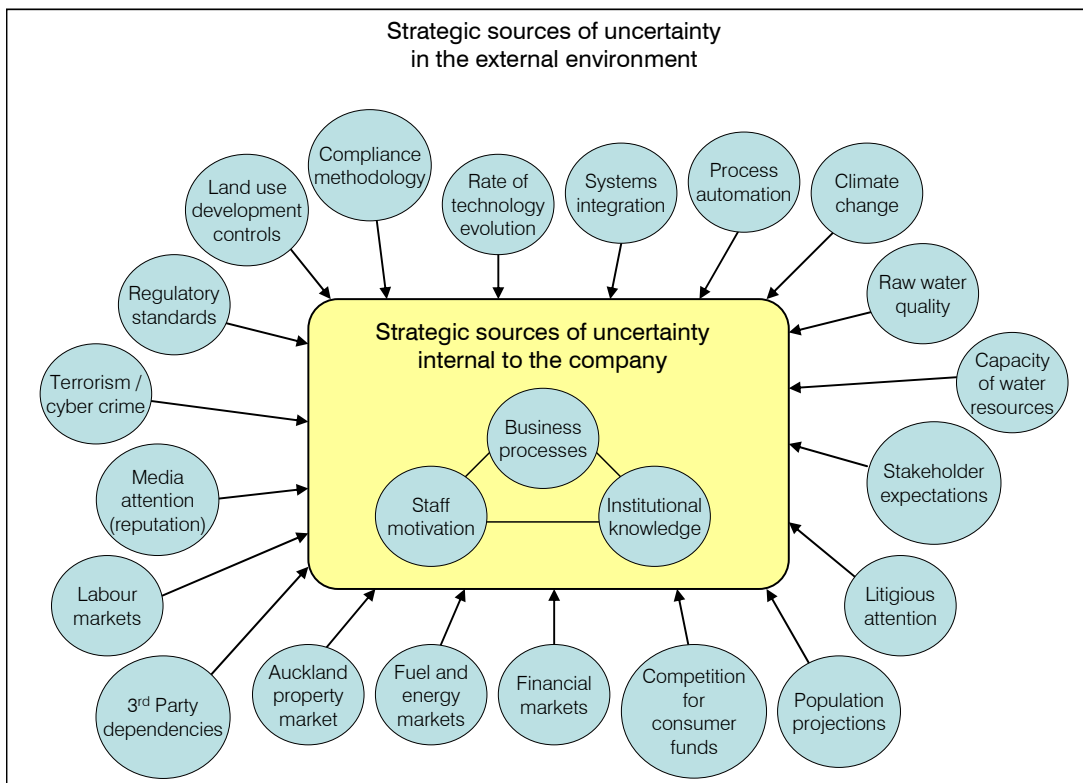
So some of the groupings that seem to stand out are around procurement, so trends in the costs and scarcity of materials, consumables, and infrastructure services, leading to heightened supply chain risk; around human resources, so concerns about the aging workforce and skilled labour markets, increasing labour costs, and problems securing the right people. Another one is concerns about the increasing complexity of technology, and also resource planning, so concerns about the future development of Auckland, and about maintaining the quality of raw water supplies. A big one related to that is of course climate change. There were quite a number of climate change related issues there. I think these kinds

---

<sup>39</sup> External: Political, Economic, Socio-Cultural, Technological, Environmental, Legal/Regulatory.  
Internal: Business Structure, Personnel, Process, Technology, Infrastructure.

of groupings imply certain sets of focal questions that could be addressed..  
(Dialogue 10, p 308)

A total of 20 external and three internal sources of uncertainty were specified in this manner. In contrast to operational or project-level risks for which a specific event with clearly defined causes and consequences could be described, each strategic “risk” or “source of uncertainty” was really just a label for a category of inter-related environmental factors and variables about which there was some uncertainty as to future trends. So, for instance, one general source of uncertainty was that around potential changes in climate patterns. This category was given the general title “Climate change” and encompassed various climate-related factors which were identified as being potentially important to Watercare. These included, among others, altered sea levels impacting wastewater discharge facilities, atmospheric warming impacting raw water quality and biological processes, and altered rainfall patterns impacting the reliability of existing raw water sources. The factors in each category reflected perceived uncertainties around both future trends in business conditions (internal or external) and the potential impacts of those trends on Watercare’s systems and processes. Each “strategic source of uncertainty” thus afforded the possibility of describing multiple future scenarios or worlds, defined by different combinations of factors, future trends, and expectations about potential impacts for Watercare.



*Figure VI.2. General systems framework for categorising strategic sources of uncertainty*

## **Fleshing out the details of Watercare's strategic risks**

The product of our one-day workshop in January 2008 was a simple framework for organising Watercare's strategic risks. However, to this point we had been working from the initial list of 330 issues identified by Watercare's managers, which, as I noted earlier, was little more than a list of relatively vague statements. While the strategic risk framework was beginning to take shape, in the sense that we had collated those initial statements into common categories, it still lacked any meaningful detail (i.e. linking perceived changes in strategic factors to potential impacts on Watercare). It was clear that the CRM would have to re-engage with Watercare's managers to extract, to the extent possible, their knowledge of those details. To facilitate this process, the CRM undertook to flesh out, in a hypothetical sense, the details for each source of strategic uncertainty using the framework described on page 116. This amounted to describing the factors and trends relevant to each source of uncertainty (drawing on the original list of issues as a basis), and extrapolating the potential impacts on Watercare. This was essentially an imaginative exercise, albeit grounded in the CRM's extant knowledge of the organisation and its systems. With this initial "fleshing out" completed, the CRM distributed the resulting draft strategic risk report to Watercare's managers prior to their attendance at a joint workshop.

The workshop was intended to achieve two purposes. First, since the framework depicted by Figure VI.2 represented a considerable amount of analytical and synthetic work by the CRM, the workshop provided an important opportunity to explain that framework to Watercare's managers, to obtain their feedback, and to address their questions. The focus here was to explain and discuss the overarching structure of the framework and the implied definition of "strategic risk", which constituted an important shift from the asset-focussed definition in use at the time. The second purpose of the workshop was to encourage the participant managers to continue the "fleshing-out" process that the CRM had begun. With the aid of a Powerpoint® template, the managers were asked to verify, correct, and expand on the CRM's hypothetical descriptions of each strategic risk, to provide a cursory evaluation of the magnitude of the risks, and to suggest possible strategic development initiatives to mitigate the risks. The CRM performed the role of facilitator, prompting and leading the discussion, while I performed the role of scribe, recording the managers' comments on each risk as the workshop progressed. The workshop provided an important opportunity for Watercare's managers to reflect on and discuss the details of the identified strategic issues within a structured framework. The framework itself, represented by the template described on page 116 and the simplified Powerpoint® template used during the workshop, proved



valuable as a means to structure and focus the discussion, and thus to elucidate the details that the managers had been unable to express on the earlier PESTLE template. The feedback obtained from the workshop provided the CRM with the necessary level of detail to complete his strategic risk report, which would subsequently serve as a key input to the strategic planning workshop to be held later in the year.

Subsequent to the workshop, the CRM undertook a survey of Watercare's managers to establish their opinions as to (a) the relative importance of the identified strategic risks to the organisation, and (b) the extent and quality of organisational knowledge about those risks. To assess the former, the CRM asked the managers to individually list the top thirteen strategic risks in order of importance, taking into account the significance of the potential impacts on Watercare's future performance, the ability of the organisation to respond to and control the risks, and the degree to which the company's exposure to the risks was controllable (Clement 2008e). The responses from the individual managers were then aggregated to establish an overall strategic risk ranking (using a simple High-Medium-Low matrix). To assess the standard of organisational knowledge the CRM, the CRM developed what he referred to as the "Knowledge Maturity" concept. The basic idea "is to gradually improve the quality of the risk evaluations and to question the level of knowledge (uncertainty) behind each. Accordingly, where required, risk evaluations are improved in three stages, from being 'guesstimated', to 'estimated', to being 'known'" (Clement 2008e, p 17). The CRM asked Watercare's managers to rank the standard of organisational knowledge about the likelihood and consequences of each strategic risk as either "High (accurate knowledge)", "Medium", or "Low (High uncertainty)", and then aggregated the individual responses to produce an overall ranking. This directly addressed the CRM's concern that risk evaluations are not informative unless the quality of the underlying risk assessment is also made explicit<sup>40</sup>. That is, the CRM recognised that while Watercare's managers were able to identify a range of strategic issues, to elaborate certain details of those risks, and to express an opinion as to the significance of those risks for the organisation, without further investigation and documentation it remained impossible to judge whether that information was based on imagination and guesswork or a detailed knowledge of the issues established through rigorous investigation and long experience. The CRM's survey sought to make this standard explicit, such that the results would be "useful for guiding and prioritising research and investigation activities within Watercare" (Clement 2008e, p 18).

---

<sup>40</sup> The CRM had identified this as an issue during his initial review of Watercare's risk management capabilities.



# Appendix VII

## Examining the CRM's contradiction

In Dialogue 12 the Corporate Risk Manager expressed a view of Watercare as a mature organisation which consistently performed well in its local context, and which arguably was “pretty good at managing the common risks associated with the business” (p 319). This perception was supported by consideration of the firm’s institutional and operating contexts, and long-term enterprise performance.

The entity known as Watercare was formally constituted in 1992, but the organisation to which that name was given was itself much older, having existed in various forms for more than half a century. While organisation age does not necessarily correlate with capability maturity or performance, the following points may be noted (see Appendix I for a full overview of the Watercare enterprise and its performance record):

- The fundamental mission of the organisation, to deliver bulk water and wastewater services to the Auckland Region, had not changed. Watercare’s overarching governance structure was also static, having remained unchanged since 1992.
- Watercare’s physical infrastructure was strongly path-dependent due to multi-year to multi-decadal rates of renewal. This was evident in both the technological form and geographical extent of Auckland’s water and wastewater networks, reflecting the history of the systems dating back to the construction of the first dams in the Waitakere Ranges in the early 1900s and the opening of the Mangere wastewater treatment plant in 1960 (Murdoch 1993; Watercare Services Ltd. 2003b).

- The dominant drivers of Watercare’s investment programme were the underlying trends of urban development, land-use change, and population growth across the Auckland region, regulatory changes affecting the company’s levels of service, and the idiosyncratic needs and demands of the LNOs. None of these drivers were highly dynamic, at least relative to the capabilities of the Asset Management Group to adapt and respond to any such change. The company regularly scored well on Asset Management benchmark tests, and was considered to be a best practice organisation by its industry peers.
- While the company’s performance record was not completely untarnished, it consistently performed well against key performance indicators, and was arguably effective with respect to fulfilling its core legislative objective to balance the long-term costs of maintaining service delivery to required standards against the long-term risks to the company’s assets.

In this context Watercare was arguably a mature, high-reliability organisation with well-established and well-tested practices for planning, developing, and operating water and wastewater infrastructure. This view (referred to here as View 2) can be seen as contradictory with respect to the CRM’s earlier observations about Watercare’s “risk data” and “risk culture” in light of his assumed decision support rationale (referred to here as View 1). The apparent contradiction is made explicit in the following sections and in Table VII.1 (indicated by the arrows).

*Table VII.1. Contradiction between the CRM’s expressed views of the Watercare organisation*

	<b>View 1</b>		<b>View 2</b>
<b>Actual empirical observation</b>	That decision making at Watercare was informed by either poor quality “risk data” or no “risk data” at all.		That Watercare consistently performed well against all KPIs.
<b>CRM’s Rationale</b>	The quality of “risk data” was positively correlated with the quality of organisational decision making and, therefore, with the performance of the organisation as a whole		The quality of “risk data” was positively correlated with the quality of organisational decision making and, therefore, with the performance of the organisation as a whole
<b>Expected empirical observation</b>	That Watercare’s performance should have been consistently poor or erratic against KPIs.		That decision making at Watercare should have been consistently informed by good quality “risk data”

## The implications of View 1

The CRM had valid concerns about the quality of Watercare’s “risk data” (described in Chapter 3). Risk was a key factor in all capital decisions, and the CRM was concerned that risk evaluations performed under the 2003 framework were producing risk profiles (i.e. “risk data”) bearing only an unreliable and inaccurate relation to the real probability of future impacts on the company’s performance. The framework did not objectively reflect the company’s actual corporate objectives and performance standards, and therefore implied a significant degree of subjectivity in the assessment of risks. There was evidence that risk evaluation and communication were being adversely affected by that subjectivity. The risk entries in Watercare’s risk register were evidently incomplete, ambiguous, and subjective, as well as being non-transparent as to their sources. And the graphical representations of Watercare’s risk profiles were mathematically problematic. These issues implied problems with the quality of staff knowledge about the company’s risk exposures more generally; a conclusion that was apparently supported by the results of the CRM’s staff survey (analysed in Chapter 3), which led him to conclude that a significant percentage of staff did not use the risk register and were not familiar with the risks associated with their roles and responsibilities.

The CRM’s decision support rationale (made explicit in Chapter 3) was rooted in an assumption that good “risk data” was positively associated with the quality of organisational decision making and, therefore, with the performance of the organisation as a whole; where good “risk data” was supposed to be that which objectively represented the systems, processes, and environments about which, and within which, decisions are made. The CRM’s observations about Watercare’s “risk data” and “risk culture” suggested that decision making at Watercare was generally informed by either poor quality “risk data” or no “risk data” at all. Combining these observations with the CRM’s decision support rationale leads to an implication of generally poor organisational decision making and enterprise performance (see the left-hand side of Table VII.1). But this inference contradicts the empirical evidence that Watercare was a mature, high-reliability organisation which consistently performed well against key performance objectives.

## The implications of View 2

Performance is always relative to the task at hand, and, consequently, to the context in which that task is performed. If an actor performs a task well (or poorly) in a particular instance,

then that success (or failure) may or may not have been the actor's fault. But if that actor consistently performs well (or poorly) over a long enough period of time then a judgement may be formed as to the competency of that actor to perform that task (Ryle 1949; Tsoukas and Vladimirou 2001). It does not matter whether the actor is able to perform that task only in one context but not in others, it is still possible to form an opinion as to the competency of that actor with that task in that context. Importantly a judgement of competency implies a judgement as to the quality of decision making involved in performing that competency. That is, a competent actor is one who generally makes "good" or at least adequate decisions with respect to task and context (without necessarily implying that they always do; see Chapter 6). Vice versa, an incompetent actor is one who generally makes "poor" decisions with respect to task and context (without implying that they always do).

This line of reasoning may, I think, quite reasonably be transposed to the organisation as a whole. If the organisation is observed to perform consistently well (or poorly) over a long enough period of time, then a judgement may be formed as to the competency of the organisation (Tsoukas and Vladimirou 2001). This implies a judgement as to the quality of organisational decision making. An organisation that consistently performs well is one that generally makes "good" or at least adequate decisions (without necessarily implying that all decisions made by that organisation are "good"). The judgement is relative to the activities of the company and its context of operations, which negates the need for an absolute judgement of decision making quality. That is, the judgement does not imply anything about the particular characteristics of that decision making which conferred "goodness", but simply that the decision making was "good", or at least adequate, on a general basis.

On this basis then, the observation that Watercare consistently performed well against key performance indicators leads to the conclusion that decision making within the Watercare organisation was also generally "good" (without implying that it was always so). If the CRM's decision support rationale is assumed to hold in reverse, then one would expect to find that the company's "risk data" was also generally "good", which, of course, contradicts the evidence to the contrary (see right-hand side of Table VII.1).

## **Resolving the contradiction**

The disagreement between observation and expectation in Table VII.1 implies problems with one or more links in the chains of reasoning leading from the empirical observations to the corresponding predictions. That is, the contradiction may be explained away if one or more of the following are at least partially true:

- i) The CRM's concerns about the quality of "risk data", the objectivity of the 2003 risk framework, and the state of Watercare's risk culture were unjustified; and/or
- ii) The review of Watercare's KPIs in Appendix I was not sufficiently accurate or comprehensive to reveal the true nature of the company's performance; and/or
- iii) The assumed correlation between the quality of "risk data" and the quality of organisational decision making was incorrect; and/or
- iv) The assumed correlation between the quality of organisational decision making and enterprise performance was incorrect.

The empirical observations, (i) and (ii) above, have been justified at some point in this thesis. The accuracy and extent of my review of Watercare's performance was justified in Appendix I, while the CRM's concerns about Watercare's "risk data" and "risk culture" were justified in Chapter 3. Since the latter part of the CRM's rationale, (iv) above, is a general assumption of most organisation and management theory it may also be treated as justified for the purpose of this discussion.

Of the above points, the one most likely to be true was (iii). That is, the weak link in the CRM's rationale and probable root of the above contradiction was the assumed correlation between the quality of "risk data" and decision making. The initial relationship was intuitive, represented by the CRM's expectation that "good decisions should be based on good data". In Chapter 3 this intuition was made to resist deconstruction by appeal to the following assumptions:

- That "good" decisions account for risk, or at least address uncertainties about performance via some form of risk assessment. This was revealed as the underlying value claim for Risk Management: for any decision, if the system and its performance objectives are comprehensively defined, and if uncertainties (risks) are comprehensively identified and rigorously assessed, then threats are more likely to be mitigated and opportunities are more likely to be realised, leading to improved and more reliable performance, compared with implicit, reactive, and unplanned approaches to management.
- That the better the data or information at one's disposal the more informed one should be about the decision at hand, and, all else being equal, a more informed actor should make better decisions. The assumption was that the quality of "risk data" was positively correlated with knowledge about risk and hence with the quality of decision

making. It was argued that when the CRM talked about “good data” he meant information that was objective, in both senses of the term, i.e. information characterised by both correspondence to reality, and the constraint of personal or political bias.

Chapter 3 described how the CRM’s vision for better “risk data” was expressed in his plan to redesign Watercare’s risk frameworks and registers, and to promote analytical tools and procedures for the calculation of risk. In particular, the CRM envisaged that risk should be quantified as a probability distribution, such that individual risks could be combined via probabilistic modelling to produce aggregate system risk profiles. Of course, numbers may or may not be any more objective than qualitative representations of risk. Watercare already employed numerical representations of risk, but these were, according to the CRM, mathematically meaningless. Rather, what the CRM envisaged was a future where the company would possess the necessary analytical capabilities, and hence objective basis, to properly quantify risk profiles within and across the enterprise. In this regard, the CRM’s notion of “good” versus “poor” data reflected the original Knightian distinction<sup>41</sup> between Risk (where probabilities are available to guide choice; i.e. “good risk data”) and Uncertainty (where information is too imprecise to be summarised by probabilities; i.e. “poor risk data”).

Both of the above points were reasonable normative assumptions about the general relationship between data, knowledge, and decision making. It could be suggested, however, that the root of the contradiction in Table VII.1 lay not with the general validity of these assumptions, but rather in the expectation that they held for Watercare specifically. In other words, Watercare’s “risk data” perhaps did not correlate well with what actors actually knew about risks in particular decision making contexts within the organisation. Such an explanation would, in fact, be consistent with the CRM’s concerns about the company’s existing (2003) risk assessment framework (described in Chapter 3). The CRM’s primary concern about that framework was that it did not objectively represent the company’s corporate objectives or performance standards or parameters, and therefore did not allow actors to communicate their knowledge of risk objectively. In order for Watercare’s staff to communicate their operational knowledge to management, they had to translate that specific knowledge into the unfamiliar and arbitrary terms of the risk framework. This suggested that the “risk data” contained in Watercare’s corporate risk register and presented in capital expenditure requests was, consequently, only a poor quality facsimile of actors’

---

<sup>41</sup> See the discussion section of Chapter 4.



actual knowledge of the represented risks. The contradiction in Table VII.1 might therefore be explained if organisational actors generally possessed good knowledge of and took account of risks in their decision making, but this knowledge was not objectively represented by the so called “risk data” in the corporate risk register.

Such an explanation would, however, raise several questions: If actors possessed objective knowledge of risks, then why didn't they record this in the register? If they knew enough to be able to properly define and quantify risks, then why would they resort to such vague and ambiguous terminology in the register? Or, if the good risk information was not stored in the corporate risk register, but elsewhere, then what purpose did the poor quality copy in the register serve? Why record it at all? And how is it that this situation could persist given that it seems contradictory in its own right?

While the CRM's concerns about the effect of Watercare's existing (2003) risk framework on the representation and communication of risk with the company were reasonable in theory, it seems unlikely that the company's “risk data” was entirely disconnected from what actors actually knew about the risks they were representing. Rather, the fact is simply that most risk assessments at Watercare were performed in contexts where there was either no reasonable basis for precisely quantifying the probability distributions of outcomes, or if such a basis did exist, the cost and difficulty of the calculation made it impractical.

This was epitomised by the risk assessment performed for the proposed Hunua No. 4 bypass of Watercare's Redoubt Road reservoir complex. The bypass was intended to mitigate the effects of the potential loss or compromise of the reservoir complex by providing a direct feed from the Ardmore treatment plant to the city. The bypass was expected to add between NZ\$30 – 100 million to the cost of the HN4 project, and there was considerable uncertainty about how much security it would actually provide to Auckland's water supply. The bypass might have been anywhere from 10,000 times more reliable to no more reliable than the No Bypass option depending on the relative failure probabilities of the various assets. The substantial cost of the bypass, the limited number of scenarios for which it would actually be effective, and the relatively low likelihoods of those situations, demanded an objective, quantitative answer to the question of to what degree the bypass would actually improve the reliability of Watercare's water supply network. That answer required, however, an analysis of the reliability of Watercare's Ardmore water treatment plant under substantially abnormal operating conditions (where the plant would have to operate in the absence of the volumetric buffer normally provided by the Redoubt Rd reservoirs). Those conditions rendered the existing reliability models of the plant inapplicable (since they modelled normal operating conditions), such that the only remaining basis on which to calculate the plant's reliability

under those conditions was to appeal to the personal knowledge and experience of the plant manager. His estimate that under those conditions a major failure could be expected within a few weeks may well have been accurate, but the point is that the situation precluded the quantification of a precise probability distribution of outcomes.

The Hunua No. 4 Redoubt Road Reservoir Bypass assessment was typical of risk assessment contexts at Watercare. Chapter 4 described how the CRM guided Watercare's managers through a process of identifying and assessing the company's strategic risks. Although the CRM eventually succeeded in resolving the managers' initially vague perceptions of environmental uncertainty into 23 defined strategic risks, those "risks" described large-scale, long-term, and multi-variable changes in the firm's internal and external environments which would not have been amenable to quantification with reasonable certainty (and certainly not without great cost). Thus, despite the potential significance of such strategic changes, objective quantitative estimates of those risks were beyond the calculative abilities and resources of the firm.

The objective quantification of project risks was similarly difficult. As one project manager commented in Dialogue 5 (p 281) some processes and activities were common to all projects, but all projects were also unique. This feature alone precluded objective quantification of probability distributions for risk outcomes.

A large proportion of Watercare's operational risks concerned the potential failure of elements of the physical infrastructure networks. A reasonable basis did exist for the quantification of a sub-set of these risks in the form of empirical information on the condition and reliability of certain assets. This information formed the basis for Watercare's reliability modelling programme, the purpose of which was to quantitatively model the reliability of the company's water and wastewater treatment plants and pumping stations. To my knowledge this was the only time and place where an attempt was made to objectively quantify probability distributions for risk outcomes. For the remainder of Watercare's physical assets, however, no such information existed, or if it did it was insufficient. It was this fact which thwarted the CRM's plan to extend the company's reliability modelling capabilities in Chapter 4.

Further, and somewhat ironically, when actors were presented with the means to specify risk in more objective terms, they insisted on comparing the resulting data against existing outcomes, i.e. against the very data about which the CRM and I were concerned. Appendix II, for instance, describes my efforts, in 2005, to redesign the Operations risk assessment framework with the help of Watercare's Operations Risk Specialist. Although the resulting assessment template allowed users to specify risks in terms specific to the Operations

context, and was therefore (at least in our opinion) a substantial improvement over the existing framework, its acceptance by Watercare's engineers as a viable tool was contingent on a comparison of the outputs from the new template against the existing assessments contained in Watercare's risk register. Similarly, in Dialogue 2, when the CRM described his proposal for an analytical procedure for prioritising capital projects using reliability modelling, Watercare's Principal Water Planner responded that the outputs of such a method would have to be calibrated against "the outcomes that we have now" (p 249). In both of these cases, the apparently subjective and problematic status quo was held up as the baseline for evaluating the appropriateness of new and supposedly more objective methods of assessing risk. A further, similar example is described in Dialogue 14. In that dialogue the CRM referred to a situation where one of Watercare's project managers added what she referred to as a "Sanity Check" sheet to the CRM's new project risk register (p 333). Once the risk assessment had been completed this new sheet sorted the identified risks in order of magnitude, the purpose of which was to allow the project manager to check whether the resulting risk ranking "looked about right". In other words, the supposedly more objective outputs of the new risk assessment tool were checked against the project manager's gut feeling as to the relative priorities of the identified risks.

Thus, it would seem that Watercare's "risk data" was consistent with how actors within the company actually understood or perceived risk. Indeed, if it were not then Watercare's managers and staff would not have explicitly referred to the register as the repository of information on the company's identified risks, nor referred to graphical representations in reports and Capex requests as the primary means of communicating the significance of those risks. While that "risk data" may have been sub-standard, as judged against certain abstract expectations of what risk should look like, it was apparently not problematic as far as organisational decision making was concerned. That is, Watercare's managers and employees may not have been able to specify and quantify risk in terms consistent with the theoretical definition (i.e. as a quantified probability distribution), but this did not stop them making good or at least adequate decisions, as indicated by the long-term performance of the enterprise. If this is accepted then the root of the above contradiction lay not with the rationale that more informed decision makers generally make better decisions, but rather with the implicit assumption that the reverse is also true, i.e. that "good" decision making requires a certain quality of information or knowledge. The above contradiction is therefore resolved by the realisation that what Watercare's management and staff knew about the company's risks was appropriate to their respective decision making purposes, at least most of the time.



# Appendix VIII

## Decision: cutting off the framing process

Rationality describes the relationship between knowledge, belief, and action, and a decision is the activity which both embodies and constitutes that relationship. This appendix supports the discussion in Chapter 6 by exploring the question of what it means to move rationally from knowing to acting.

### **On moving from knowing to acting**

It is a basic assumption of pragmatism that the truth of a matter can never be proven in an absolute sense due to the epistemic limits which confront human inquiry (see Chapter 2). These limits mean that the extent to which the human mind can get an adequate grip on reality is unavoidably limited; that all human knowledge is to some extent imperfect, and, in principle, fallible; and that humans can never aspire to perfect knowledge, but only to improve on existing knowledge (Loasby 1999, pp 1-7; Rescher 1995, pp 38 & 50; Rosa 1998, p 34). Faced with these limits, one can never prove, in an evidentiary sense, the absolute nature of reality, nor the absolute completeness of a system of beliefs, nor the practical value of a proposition for all time (Rescher 1995, 2005). While this does not stop people from believing all manner of things, it does throw into question the notion of what it means to know something, and how we achieve the movement to action in the face of incomplete knowledge. The sections below address the foundational questions of:

- How does belief operate?
- Is truth necessary for rational action?
- What constitutes justified belief?
- What constitutes reliable evidence?
- What constitutes sufficient evidence for rational belief?
- What are the normative conditions for rational action in the face of incomplete knowledge?

### **How does belief operate?**

Belief is one of the three conditions for knowledge, at least as that concept has traditionally been defined (Cambridge Dictionary of Philosophy 1999; Steup 2010). The notion of what it means to believe something is philosophically contentious, but the term “belief” is generally taken to refer to the attitude that a person has when they take something to be the case or regard it as true (Schwitzgebel 2010). I think it is sufficient for the purposes of this discussion to regard belief as involving a psychological relation – that is, the believer must, in some way, be psychologically related to that which is believed, and it is this relation that we refer to as belief (Cambridge Dictionary of Philosophy 1999; Schwitzgebel 2010). While there is no need to assume, a priori, that that relation should or must take a particular form in practice, I am assuming that belief operates in such a way that the concept of a psychological relation holds for both System 1 (intuitive) and System 2 (reflective) cognition (Kahneman 2003; Weber and Johnson 2009; see Chapter 6). That is, I am assuming here that the psychological relation between thought and action which is established via reflective calculation (System 2 cognition) is the same kind of relation as that established via intuitive calculation (System 1 cognition). This is a simplifying assumption, since if it is not the case then it would be necessary to understand how belief operates vis-à-vis System 1 and System 2 processes, and the interrelation between the two kinds of process.

In order to establish whether this assumption is reasonable, it is necessary to consider how belief manifests in practice, and especially in contexts of skillful action, since it is intuitively obvious that we do not go around affirming propositions every time we decide something. In other words, if it is true that most people, most of the time, do not explicitly affirm their belief in propositions (or their equivalents) when making decisions, then what does it mean to hold such belief?

I could appeal here, however weakly or conveniently, to the general assertion that belief

need not be active in a conscious sense, i.e. a person might believe a vast number of things but this does not mean he or she is consciously aware of all of them at any single moment in time (Cambridge Dictionary of Philosophy 1999; Schwitzgebel 2010). This is evidenced by the fact that people generally believe all manner of relatively mundane things (e.g. that they have heads, that Wellington is the capital of New Zealand, etc.) without having to consciously think those things (Schwitzgebel 2010). On this basis it would be possible for an agent to believe certain propositions without having consciously formulated them. If so, then the way in which I conceptualise rationality here would not be inconsistent with contexts of skillful action where agents calculate in a largely sub-conscious and therefore implicit manner.

However, a stronger and more problematic objection can be derived from how belief is thought to relate to “knowing-how” (Ryle 1949). For instance, in his overview of the concept of belief, Schwitzgebel (2010) notes as follows:

There may also be types of knowledge that are not types of belief, though they have received less attention from epistemologists. Ryle (1949), for example, emphasizes the distinction between knowing how to do something (e.g., ride a bicycle) and knowing that some particular proposition is true (e.g., that Paris is the capital of France). In contemporary psychology, a similar distinction is sometimes drawn between procedural knowledge and semantic, or declarative, knowledge (see Squire 1987; Schacter, Wagner, and Buckner 2000; also the entry on memory). Although knowledge-that or declarative knowledge may plausibly be a kind of belief, it is not easy to see how procedural knowledge or knowledge-how could be so, unless one holds that people have a myriad of beliefs about minute and non-obvious procedural details. At least, there is no readily apparent relation between knowledge-how and “belief-how” that runs parallel to the relation epistemologists generally accept between knowledge-that and belief-that.

This suggests that my assumption that belief operates in essentially the same way vis-à-vis both System 1 and System 2 cognition might be incorrect. The problem is this: it may be the case that the kinds of sub-conscious cognitive processes involved in calculation might fall into the same category as those involved in knowing-how, and if knowledge-how does not involve a corresponding belief-how relation, parallel to the knowledge-that/belief-that relation typically assumed in epistemology, then this would imply that calculation does not necessarily involve belief as I have assumed here.

I can, however, appeal to the case put forward by Stanley and Williamson (2001), that knowledge-how is in fact a species of knowledge-that. Ryle’s (1949) original argument against the “intellectualist legend” seems eminently sensible, and it is perhaps for this reason that his thesis that “knowing-how” is fundamentally different from “knowing-that” remains today widely accepted and influential, both within epistemology and beyond (Stanley and Williamson 2001). In their paper, however, Stanley and Williamson (2001) mounted a strong

critique of Ryle and demonstrated that both his thesis and his account of knowledge-how are false (Stanley and Williamson 2001, p 441). They then developed their own positive account of knowledge-how, justified by appeal to “well-entrenched doctrines of linguistic theory”, according to which knowledge-how is a species of knowledge-that (Stanley and Williamson 2001, p 444). Although I cannot claim to have followed the development of Stanley and Williamson’s account in detail because I am not versed in linguistic theory, their argument has one critical implication with respect to the above question. That is: if one accepts the premise for knowledge-that, that an actor can be said to believe a proposition without being occurrently aware of it, and even that an actor can retrieve a proposition from memory without being able to express it in non-indexical words, then one must also accept it for knowledge-how, if knowledge-how is a species of knowledge-that (Stanley and Williamson 2001, pp 439-440). In other words, if Stanley and Williamson are correct, then knowing-how must involve belief in the same or similar ways as belief is usually taken to relate to knowledge-that; which would deny the above objection.

The above points also address the objection that my conceptualisation of the decision making process implies an infinite regress with respect to an agent’s beliefs. That is, does belief in a given proposition, *p*, also imply belief in all antecedent or conditional propositions? It follows from above, and indeed from common sense, that an agent need not actually affirm any such antecedent beliefs in a conscious manner, even if he or she consciously affirmed belief in *p*, nor for that matter even cognitively process antecedent beliefs within some deep sub-conscious area of the brain.

I might also avoid the above objections by framing the condition that rational choice requires the agent to justifiably believe certain propositions as a normative condition of making a claim to rationality after the fact. The distinction here is between the claim of “I am a rational person”, which is a claim about one’s reasoning skills in general (which we might call *rationality as intent*), and the claim of “I decided rationally”, as a claim about one’s performance in a particular decision making context (which we might call *rationality as performed*). In order to explicitly claim, *after the fact*, that a choice *was* made rationally, an actor would have to consciously, and therefore explicitly, acknowledge that he or she believed certain propositions to be true at the time the decision was made. While there are, of course, certain methodological difficulties with such post-hoc rationalisation, it is also true that in a great many contexts it is simply impossible for actors to engage in such explicit affirmation of their rationality prior to the moment of decision (e.g. driving a race car at 200 mph), while in many other contexts such affirmation is usually left implicit (such as in planning the route of a watermain). Indeed, we humans rarely reflect on our own rationality unless something has



gone wrong, which points up the fact that such explicit acknowledgement would only have to be made in order to demonstrate and hence establish a claim to rationality under conditions of audit in a time and place removed from the original decision.

### **Is truth necessary for rational action?**

A second condition of knowledge is that of truth, because one can not claim to know something which is false (Steup 2010). However, as far as reason is concerned, the objective truth of a proposition (statement, argument, or claim) is irrelevant; what matters is only that one must believe it to be true at the moment of decision. The implication is that one does not need to “know” something, in the traditional sense, in order to reasonably decide and act; and also the corollary, that one can reasonably (though probably not reliably) decide and act on the basis of information later shown to be false. I take this position, that it is justification rather than objective truth which matters as far as reason is concerned, in order to be consistent with the later criticism of the conventional notion of rationality. That is, if true belief were a requirement of reason then decisions performed in complicated situations (i.e. domains of high uncertainty and novelty) would, by definition, be less than rational. Such a definition would preclude the possibility that an agent could decide rationally in the absence of complete knowledge; and would, on this basis, be perverse.

### **What constitutes justified belief?**

The third traditional condition of knowledge is that of justification, referring in a general sense to the basis on which one believes something (Cambridge Dictionary of Philosophy 1999; Steup 2010). It is a necessary condition of knowledge because one cannot claim to have known something if the belief in question only turns out to have been true by accident or luck (Steup 2010). Although I have just argued above that truth is not a condition of reason, it still follows that one can not reasonably claim to believe something to be true without a reason, else the belief would be arbitrary.

In philosophical debate controversies surround both the meaning of the term “justification” and the question of what makes beliefs justified, i.e. the substantive conditions for justification (Cambridge Dictionary of Philosophy 1999; Steup 2010). In light of the earlier pragmatic assumption about the limits of human abilities to know the world, a definition of justification as evidence will suffice here. That is, an agent is justified in believing that a proposition,  $p$ , is true if the agent has adequate indication, in the absence of overriding evidence to the contrary, that  $p$  is true (Cambridge Dictionary of Philosophy 1999). This is

consistent with the concepts of partial belief and plausible reasoning, i.e. the possibility of being more or less confident about the truth of a proposition based on the kinds of evidence which support it (Jaynes 2003). Furthermore, a broad understanding of justification as evidence is, at least for the purposes of this project, not incompatible with normative theories of justification since normative reasoning as to the morality of actions and outcomes can constitute evidence for or against the truth of a proposition, but is necessarily subject to the same epistemic limits as all other knowledge.

Debate over what makes beliefs justified concerns questions of the nature and reliability of evidence. That is, what constitutes evidence and on what basis may the quality or strength of evidence be judged? Evidentialists argue that what matters for justification is the possession of experiential evidence, such that an agent is justified in believing that  $p$  if the agent has an experience which represents  $p$  as being true; while reliabilists argue that evidence alone is not sufficient and that the source of the evidence must also be reliable (Steup 2010). The second element of justification is the question of sufficiency. That is, if we can never prove truth in an absolute sense, but we do not anyway need absolute proof of truth in order to form and hold beliefs, and if, for any given proposition, we have access to various kinds and strengths of evidence, both for and against the truth of that proposition, then what constitutes sufficient justification or “adequate indication” of the truth or otherwise of that proposition? The following sections present the positions adopted here with respect to the questions of the reliability and sufficiency of evidence.

### **What constitutes reliable evidence?**

*The reader is referred to Chapter 2 which addressed the question of what constitutes reliable knowledge in the pragmatist perspective.*

### **What constitutes sufficient evidence for rational belief?**

There is no absolute answer to the question of sufficiency, but a minimum condition may be stipulated, at least in a general sense, by drawing on notions of partial belief and plausible reasoning (Jaynes 2003; Ramsey 1926). Taking into account the various kinds and strengths of evidence, and notwithstanding the practical ability of agents to objectively perceive, understand, and evaluate that evidence, belief in a proposition may be justified when the evidence for the truth of that proposition outweighs the evidence against, or outweighs evidence for the truth of competing propositions. That is, all that is required to rationally form a belief is weight of evidence. There may be considerable evidence to the contrary, but

as long as the evidence in favour outweighs the evidence against, however fractionally, then a belief in the *probable* truth of the proposition can be justified. This provides the minimum condition for justified belief on which basis an agent might proceed with subsequent cognitive or practical endeavours. However, some clarifications must be noted.

First, if one can hold a justified partial belief, if one can reasonably believe the truth of a proposition by degree, then this necessarily implies that one believes the proposition is more likely to be true than false (Ramsey 1926). This is equivalent to saying that the proposition,  $p$ , is probably true, and can be represented as:  $0.5 < P(p) < 1$  (Ramsey 1926). Where a proposition concerns the expected outcomes of future actions, as in “action  $x$  will lead to outcome  $y$ ”, justified belief requires the agent to believe that outcome  $y$  will be more likely than the cumulative likelihood of all other possible outcomes of action  $x$ . If this is not the case, then the agent cannot claim to believe that the proposition is probably true, even if outcome  $y$  is the most likely of the various possibilities.

Second, the strength of one’s belief in the truth of any proposition depends on the type and quality of the evidence which supports that belief, and that belief may only be justified, minimally or otherwise, on a *full* accounting of all the evidence at hand (i.e. the evidence which is perceived by and available to the agent). That is, to the extent that one is aware of evidence to the contrary, one cannot rationally discount or ignore that evidence in forming one’s belief, unless there is good reason to consider that evidence as false or unreliable. To do so, would be to form one’s beliefs in an arbitrary manner. It follows that improving the strength of one’s belief in the probable truth of a proposition requires one to acquire more or better evidence, which in practical terms implies a trade-off, since the acquisition of more or better evidence always involves a cost in time and resources. That is, at some point the benefit of having more or better evidence, in terms of the improved reliability of one’s beliefs, will be outweighed by the cost of acquiring that evidence; and there remains the possibility that one may not be able to acquire further evidence at all due to the limitations of time and resources, and the existing state of technology. The subsequent and fundamental question of how agents may reasonably proceed under these limitations is the subject of the second part of this chapter.

Third, the above assumptions clearly reflect the Bayesian perspective that degrees of confidence are ultimately subjective (Hansson 1994; Jaynes 2003). In this view, beliefs may differ from agent to agent, even when those agents are presented with the same information, and an agent’s degree of belief may change when the agent is exposed to new evidence (Hansson 1994). However, I do not see that it is subsequently necessary to assume that probabilities must also be taken to be purely mental phenomena (Hansson 1994). Rather, it

seems far more practical to assume that probabilities can represent the objective frequencies of phenomena in the world, or the subjective degrees of confidence of the Bayesian subject, or both, depending on the situation, and, further, that the two types of probability are interchangeable and combinable for calculative purposes; while recognising that they imply very different forms of measurement, and may take on different forms when quantified (Hansson 1994; Jaynes 2003). This is not to say that one can arbitrarily substitute subjective degrees of belief for objective estimates of probability, but rather that (i) estimates of objective probability, in the frequentist sense, are claims about the world about which an agent might be more or less confident depending on the evidential basis for those estimates (cf. the discussion in Chapter 4); and (ii) in the absence of estimates of objective probability, subjective probabilities in the Bayesian sense can provide a reliable basis for action (Jaynes 2003).

### **What are the normative conditions for rational action in the face of incomplete knowledge?**

Due to the epistemic limits which confront human inquiry, there is always a gap between our knowledge and the objective reality it represents. In this regard, uncertainty can never be completely eliminated (the world can never be framed absolutely). Much of the time, this uncertainty is either irrelevant or trivial with respect to action; we are either unaware of it or we ignore it in our calculations. But in some cases we may be required to act in situations of incomplete knowledge, where we are in some significant sense, less than completely confident in the truth or objectivity of the knowledge at hand. In other words we may be required to act on partial belief.

While partial belief may reasonably be considered justified in the minimal sense described above, such partial belief would not necessarily provide a reliable basis for subsequent endeavours, whether practical or cognitive (i.e. for action or for belief in subsequent propositions). Consider, for instance, the proposition,  $p$ , that action  $x$  will lead to outcome  $y$ . This proposition is implied in all decision making situations involving expectations about the outcomes of future actions. The minimum condition for justified belief, above, corresponds to a situation in which the weight of evidence is slightly in favour of the probable truth of this proposition. For the sake of example this might be represented as:  $P(p \text{ is true}) = 0.6 \mid P(p \text{ is false}) = 0.4$ , with  $P$  being probability in the Bayesian sense of the odds that an agent might assign, on consideration of the available evidence, to the likelihood of the outcomes  $y$  or *not*  $y$  occurring (i.e. if he were a betting man). But while the agent may be justified in

believing the probable truth of the above proposition, his own subjective expectation is that he would only be successful in his subsequent endeavours slightly more than half the time. That is, if action  $x$  were a repeatable exercise then the agent would expect outcome  $y$  to occur only slightly more often than all the other potential outcomes (i.e. *not*  $y$ ). In this regard, the agent's belief might be justified but it would not be a particularly reliable basis for action. It would clearly be advantageous for the agent to have a greater weight or strength of evidence such that he could have greater confidence in the probable truth of the proposition.

When calculating how to act in situations of incomplete knowledge the decision maker must account for the probability of outcomes other than those desired or expected. The question of how to properly account for uncertainty is the subject of traditional decision theory (see for instance Hansson 1994 for a concise overview). The purpose of traditional decision theory, both normative and descriptive, is to elaborate the rules under which agents should or do, respectively, calculate the rational choice in a given situation. Different rules apply depending on an agent's state of knowledge, typically categorised as decision making under certainty, risk, uncertainty, and ignorance (Hansson 1994; where the distinction between Risk and Uncertainty is that of Knight 1921, cf. the discussion in Chapter 4). In all situations the agent must seek to identify the best course of action (this being the rational imperative). In situations of certainty, risk, and uncertainty, this means that the agent must calculate the expected value of the various options, and then select that with the highest value; where "value" should be loosely understood as a measure of what the agent hopes to gain or achieve (Hansson 1994). Normatively speaking, the general form of the expected value calculation is modified for decision making under Risk and Uncertainty as follows:

- For decision making under Risk probability functions are assigned to outcomes because they are reliably known;
- For decision making under Uncertainty the probability functions of outcomes are not reliably known and must be modified by some measure of confidence.

In each of the above cases both the potential outcomes of a course of action, and the probability functions for those outcomes are known to some extent. But it is also possible for an agent to be ignorant with regard to these things. Traditional decision theory considers rules for deciding under ignorance in two cases: classical ignorance, where the possible outcomes are known, or at least identifiable, but where the agent has no reliable information about their probabilities, and the more extreme case of unknown possibilities, where the agent does not have a complete list of the consequences that should be taken into account (Hansson 1994). In these situations the rational imperative still applies, but, unlike the above

cases, an *expected* value cannot be calculated because the agent has no probability information, objective or subjective, on which to formulate any such expectation. Rather, under conditions of ignorance, a rational course of action may be calculated according to various strategies such as maximising security or minimising regret; where the agent may adopt different strategies depending on his or her degree of optimism or pessimism about the future, and aversion to risk (Hansson 1994).

The limitation of the rules and methodologies elaborated by traditional decision theory is that they assume that the decision frame has already been established (i.e. they assume the frame as given) and that all that remains is the calculation. That is, in order to perform a decision in accordance with the normative rules of traditional decision theory, the agent must already know what is to be taken into account and what is to be left out, and be able to identify which rules and methodologies are appropriate for his or her state of knowledge. However, the discussion in Chapters 2 and 6 have called attention to the performative aspects of knowing and deciding in the world. That is, decisions, like everything else, must be constructed and this is not necessarily an easy or entirely objective process, and, further, is one in which humans face significant constraints. For this reason, the next section addresses the fundamental question of how agents may reasonably decide when to stop looking for information and options and move on to the calculation. That is, how do agents reasonably decide that they know enough to act?

## **Conditions for settling decision frames reasonably**

If every decision implies a certain framing of the world, and if a frame is the product of a framing process (i.e. such that it must be seen as evolving or becoming over time), and if the world (or a particular problem situation) can never be completely framed in an absolute sense, then every decision necessarily implies a prior judgement that the frame is *sufficiently* complete that the actor may “cut off” the search for options and information, and move on to the decision. This is the concept of satisficing (Simon 1955, 1957), more broadly stated. But making that cut-off involves a conundrum.

We may imagine, for instance, that an actor has framed a certain number of options for how he might proceed with respect to his objectives, and could decide the situation at this point by choosing a “best” option from amongst those already framed. If there are no limiting factors, such as time or resource constraints, which might force that actor to take the decision immediately, then he faces a conundrum. The imperative to reason, to seek out the

most appropriate course of action, requires that he should keep searching for a better option. If he does there are four possible outcomes:

- He does not succeed in identifying any other options before his time and resources are exhausted;
- He does succeed in identifying more options, but none are better (i.e. more efficient and effective) than the best option already identified;
- He does succeed in identifying a better option, but the marginal improvement over the best option already identified is outweighed by the expense of the search;
- He does succeed in identifying a better option, where the marginal improvement over the best option already identified does outweigh the expense of the search.

The conundrum is that only one of the above outcomes leaves our decision maker better off than he is right now, but he cannot predict with certainty which is more likely, before the fact, because the necessary information is more or less unknown to him. Thus, in the absence of any further information, on what basis can our decision maker reasonably decide whether to continue the search for information or options or to satisfice right now?

## **Resolving the satisficing conundrum**

The answer to the above conundrum lies in the very notion of what it means “to frame”. The sense in which framing has been used in this thesis is that of separating, sorting, and arranging those things which are relevant with respect to a decision from those things which are not, in much the same way that a picture frame separates the picture from the wall. The important point is that a frame, in this sense, is a boundary or demarcation between two things (or two sets of things). Crucially, it is the erection of this boundary which brings those two things into existence, along with the boundary itself (Herbst 1976). That is, two things cannot be distinguished without a boundary, and one cannot have a boundary without implying a separation or demarcation between things. In order to distinguish a picture from the wall on which it hangs there must be a boundary between them. The frame is essential to our perception (and hence existence) of both the picture and the wall as separate, identifiable entities. Without the frame we could not even conceive of the possibility of the picture on the wall; there would simply be undistinguished space.

Since we automatically perceive the physical frame just by looking at the picture on the wall, we typically take the conceptual distinction between picture and wall for granted.

Nevertheless, it is in the very act of bringing the frame into existence, whether physically or conceptually, that we *simultaneously* bring into existence the picture and the wall; or in the case of a decision, the categories of what is relevant and significant, and what is not; or, in a fundamental sense, simply what is, and what is not (Herbst 1976).

The implication is that what is inside the frame (i.e. what is relevant) and what is outside the frame (i.e. what is not relevant) are both known to some extent. Indeed, one cannot judge an entity to be significant nor relate it to something else unless one knows something about it. Framing does not require full or even substantial knowledge of the entity in question, but, at a minimum, only the perception of an entity's existence, or more accurately, merely the perception of the existence of something that is other (Chia 1994; Cooper 1986). In this regard, framing is the start and the process of knowledge, as the mere act of perceiving is how we begin to come to know the world around us (it was in this sense that Herbst (1976) conceptualised framing as the fundamental logical act). The point is that one cannot have (some) knowledge of a frame and what is inside it without also having (some) knowledge of what is outside the frame.

This provides a way out of the above conundrum because it means that actors always have some knowledge of the broader problem domain in which they are located and therefore have some kind of basis on which to judge, *before the fact*, the likely outcome of any decision to proceed with the search for options. An actor may not be able to identify and specify any further options at a particular point in time, but he may well know enough to judge whether a further search will probably (or probably not) yield further options, the potential for any such additional options to be better than those he has already identified, and how much work this might involve. In the absence of time and resource constraints the decision to continue or cut off the framing process must be based on these expectations.

### **Making the satisficing cut-off (specific case)**

My concern here is with developing an account of what it means to resolve possibilities for belief and/or action in a rational, which is to say non-arbitrary or reasonable manner. Where alternative possibilities exist, and are perceived by the agent in question, then the rational choice of one possibility over the others requires the agent to reject those other possibilities as less favourable; i.e. the agent must calculate (Callon 1998b). Where the agent does not have complete knowledge of the problem domain then he or she must also reject the possibility of continuing to search for information and options as less favourable than proceeding with the decision on the basis of what is already known (this does not rule out the possibility that the



agent may be unable to decide the situation in a non-arbitrary way, or may be forced to satisfice prematurely due to time or resource constraints). The suggestion, above, was that this satisficing choice must be based on the agent's expectations about what is inside the decision frame, what is outside the frame, and the relationship between them.

These expectations can be phrased as propositions, about the truth or falsity of which an agent must establish a certain state of belief in order to proceed. For the above case where an actor must decide on a "best" course of action for achieving a given set of objectives, three such propositions may be phrased as follows:

- **Proposition 1 (P1) – Causality:** That, for each identified option, the proposed actions will produce the predicted outcomes.
- **Proposition 2 (P2) – Fulfilment of objectives:** That, for each option, the balance of potential outcomes will weigh in favour of achieving the agent's objectives, accounting for (i) outcomes desired, (ii) outcomes undesired, and (iii) outcomes not undesired.
- **Proposition 3 (P3) – Completeness:** That there are no other possible courses of action, which might be better than the best option already identified, and which could reasonably be identified within the constraints of the decision situation (i.e. that there are no other possible options for which the cost-benefit ratio of development is positive).

For P1 and P2 there are two possible outcomes. Either the agent can justify a degree of belief with respect to the truth or falsity of P1 and P2, with respect to any particular option; or the agent cannot justify belief either way. The agent may proceed in each case as follows:

- 1) If the agent cannot establish a degree of belief either way, then this is equivalent to a state of ignorance (i.e. insufficient knowledge), at least for that agent, in which case the agent may proceed as follows depending on his state of belief in P3:
  - a. Take the decision. Whether the agent includes or excludes the option for which he or she could not establish belief with respect to P1 and/or P2 will depend on how the agent deals with ignorance.
  - b. Keep framing, i.e. continue the search for information and options.
- 2) If the agent can establish a degree of belief with respect to the truth or falsity of P1 and P2 for any particular option then the agent may choose to do either of the following depending on his state of belief in P3:

- a. If the agent believes either P1 or P2 to be false then he may:
  - i. Take the decision with the other options already identified, but excluding the option for which he believes P1 or P2 to be false; or
  - ii. Keep framing, i.e. continue the search for information and options.

Or

- b. If the agent believes both P1 and P2 to be true then he may:
  - i. Take the decision with this and the other options already identified; or
  - ii. Keep framing, i.e. continue the search for information and options.

For P3 there are also two possible outcomes. Either the agent can justify a degree of belief with respect to the truth or falsity of P3; or the agent cannot justify belief either way. The agent may proceed in each case as follows:

- 3) If the agent cannot establish a degree of belief either way, then this is equivalent to a state of ignorance (i.e. insufficient knowledge), at least for that agent, in which case the agent may choose to satisfice as follows. Either:
  - a. If time, resource, or other limits have already or will soon be encountered then it would be reasonable to cut-off of the framing process and take the decision as it stands; or
  - b. If time, resource, or other limits do not impose an immediate need to take the decision then it may not be unreasonable to continue the framing process for the time being; or
  - c. If time, resource, or other limits do not impose an immediate need to take the decision then it may not be unreasonable to take the decision as it stands, while reserving the option of reframing the decision depending on what happens (i.e. proceeding on an experimental or trial and error basis).
- 4) If the agent can establish a degree of belief with respect to the truth or falsity of P3 then the agent may satisfice as follows. Either:
  - a. If the agent believes P3 to be false then he should take the decision as it stands; or
  - b. If the agent believes P3 to be true then:
    - i. If time, resource, or other limits have already or will soon be

encountered then it may nevertheless be reasonable to cut-off of the framing process and take the decision as it stands; or

- ii. If time, resource, or other limits do not impose an immediate need to take the decision then the agent should continue the framing process for the time being.

### **Clarifications and qualifications**

A number of comments must be made at this point to explain and clarify the above logic for the satisficing decision. With reference to the earlier discussion, belief is taken to be a form of psychological relation between the agent and a proposition, where that relation may be weak or strong, i.e. belief may be partial (Cambridge Dictionary of Philosophy 1999; Ramsey 1926; Schwitzgebel 2010). An agent's state of belief, which is to say his degree of confidence in the truth of a given proposition, is assumed to depend on the nature and reliability of the evidence at the agent's disposal. That is, if an agent's state of belief with respect to the above propositions is to be non-arbitrary then this requires that what the agent takes to be known, and the grounds on which he or she makes certain judgements of relevance and significance, must be justified to some extent; i.e. there must be something which lends weight or reliability to the evidence as far as the agent's belief is concerned. There is, however, a plethora of ways in which agents might perceive and come to know the world around them, and similarly for building, storing, and transmitting knowledge about the world (see, for instance, Gardner's (1983, 1999) theory of multiple intelligences and Pepper's (1942) world hypotheses (root metaphors) for theory building). Thus, the reliability of evidence may be evaluated in a variety of ways depending on the specific context, and not all forms of evidence are equivalent, or even necessarily sufficient, with respect to justifying specific cognitive and practical endeavours in specific contexts. As such I am not concerned, here, with what constitutes evidence for belief in any specific sense. Rather, my concern is with identifying those general categories of things about which an agent must have evidence with respect to the above propositions, and with understanding the relationship between belief in the above propositions and subsequent action.

### **Basis for belief in P1 and P2**

If the agent can establish a degree of belief with respect to the truth or falsity of P1 and P2 for any particular option then, regardless of whether the agent believes P1 and P2 to be true or false, this implies that the agent has already calculated the potential outcomes with respect

to that option, accounting, presumably, for uncertainty where necessary (i.e. the possibility of outcomes other than those desired). That is, in order to establish a justifiable degree of belief in P1 and P2 with respect to any option the agent must already have calculated the expected value of that option; where the term “value” should be loosely understood as a measure of what the agent hopes to gain or achieve, and where the term “calculate” does not assume or imply anything about the particular method of calculation that might be employed by the agent. The agent’s belief in P1 and P2 is therefore the product of a means-ends assessment. The basis of that assessment is made explicit in Table VIII.1, below.

*Table VIII.1. Basis for belief in P1 and P2*

<p><b>Proposition 1:</b> That for each identified option the proposed actions will produce the predicted outcomes. Belief in P1 is based on an analysis of the nature and behaviour of the world.</p>	
<p><i>Belief in P1 is founded on the following evidence:</i></p> <p>Knowledge of the present state of the world/system.</p> <p>Expectations about the future behaviour of the world/system.</p> <p>Expectations about outcomes of future actions based on some form of causal reasoning.</p>	<p><i>Degree of belief in P1 depends on:</i></p> <p>Nature and reliability of knowledge about the present state of the world/system.</p> <p>Nature and reliability of analytical framework/methodology used to make predictions in that system.</p>
<p><b>Proposition 2:</b> A prediction that for each option the balance of potential outcomes will weigh in favour of achieving the actor’s objectives, accounting for (i) outcomes desired, (ii) outcomes undesired, and (iii) outcomes not undesired. Belief in P2 is based on an evaluation of (a) the present performance of the system, and (b) expected future outcomes, against known objectives and preferences.</p>	
<p><i>Belief in P2 is founded on the following evidence:</i></p> <p>Knowledge of objectives, preferences, and performance criteria and indicators.</p> <p>Knowledge of present system performance and predictions of expected future performance under each option.</p> <p>A methodology for relating and comparing performance outcomes against objectives and preferences.</p>	<p><i>Degree of belief in P2 depends on:</i></p> <p>Reliability of knowledge about objectives, preferences, performance criteria &amp; indicators.</p> <p>Reliability of knowledge about the present state of the world/system.</p> <p>Reliability of analytical framework/methodology used to make predictions in that system.</p> <p>Reliability of the evaluative methodology.</p>

**Basis for belief in P3**

Regardless of the agent’s state of belief regarding P1 and P2 the agent is faced with the satisficing decision, and in order to proceed must establish a state of belief with respect to P3. This is a significant point. It implies that the evaluation of P3 by the agent is a continuous

process. It is not a one-off evaluation, made at a certain time or when a certain set of conditions arise, but rather is performed continuously throughout the framing process. That is, the agent is continuously engaged in a process of evaluating whether to continue or cut off the search for information and options. These two points reflect the long standing argument that humans do not make decisions via a linear progression through an ordered series of consecutive stages (traditionally: problem identification, information collection, option development, option evaluation, implementation; Dewey 1910; Simon 1960; Brim et al. 1962). Rather, if human decision making can be divided into such stages, then they are performed in parallel rather than consecutively:

We believe that human beings cannot gather information without in some way simultaneously developing alternatives. They cannot avoid evaluating these alternatives immediately, and in doing this they are forced to decision. This is a package of operations and the succession of these packages over time constitutes the total decision making process. (excerpt from Witte 1972, p 180)

With regard to P3 it is necessary to clarify the intuitive distinction between “potential” and “actual” options. Although the very concept of a frame implies a demarcation between the categories of what is within and what is without, and it is intuitively easy to associate these with the categories of “actual” and “potential” respectively, this is a false distinction. It is false because the very act of labelling something as an “option” is itself the act of framing. In so far as the term “option” can be taken to refer to a possible course of action, the act of labelling any such course of action as an option involves judging the various constitutive elements of that course of action as both relevant and significant to the decision at hand. Thus options exist, by definition, only within the decision frame, i.e. “options” being collections of entities which have been arranged into meaningful wholes and related to the decision at hand. Beyond the decision frame there are no options, merely known but unrelated and unarranged entities – “potential” in this sense refers to the as yet undistinguished possibilities for arranging and relating (i.e. framing) those exterior entities

This is a significant point. It means that while Proposition 3 relates to the agent’s expectations about finding a further option which is better than those already identified, the agent’s state of belief with respect to P3 does not depend on any kind of assessment of things which might be called “potential options”. It follows from the points already made that if the agent were in a position to explicitly assess the benefits of a “potential option” over those “actual” options already identified, and against the costs of developing that new option, then he must already have framed and evaluated the new option, and therefore established a degree of belief with respect to P1 and P2 for that option, and also its relative priority to the

other options already identified. Since this is the same evaluation which is made for all other options, it does not constitute the basis for closing off the framing process. The satisficing decision would remain open for the agent.

Nothing about what I have just said should be taken to imply that the agent is ignorant about the broader problem domain. The last point, above, refers not to what the agent knows, or the information at his disposal, but rather to what he has and has not yet framed vis-à-vis the decision at hand – this latter action implying the additional judgements of relevance and significance. While Witte’s comment, above, applies equally to these judgements, such that upon coming to know something we automatically tend to judge its relevance and significance to the tasks at hand, it is also true that this is a process that can take considerable time, depending on the circumstances. Thus an agent might instantly relate new knowledge or information to the decision at hand, or it might take hours, days, or weeks for those distinctions to be realised. In this sense, an agent might in fact be quite well informed about the broader problem domain, and yet still not have framed further options.

This conceptualisation of the relationship between what is inside and what is outside the decision frame reflects the claim by Tsoukas and Vladimirou that “knowledge is the individual capability to draw distinctions, within a domain of action, based on an appreciation of context or theory, or both” (2001, p 973). That is, the process of framing, of bringing an entity within the decision frame, involves drawing distinctions about that entity and its relationship to the problem or decision at hand. The important concept here is that what is or has been framed is distinguished from all else that is known by the agent, and from information in general, by the judgements of relevance and significance.

Proposition 3 relates to the agent’s expectations about continuing to make such distinctions going forward, and about whether this might significantly alter the decision at hand, i.e. would such distinctions result in the agent making a significantly different choice than he or she would otherwise? The fact that the agent has not yet made those distinctions need not imply that the agent is ignorant, or even significantly uncertain about them. This may be illustrated by a commonplace scenario. When buying a car, an agent’s objectives and decision criteria are defined to the extent that he knows generally what sort of car he is looking for, and he has a general set of preferences with respect to make, colour, upholstery, and price range. Our consumer’s options include those cars that he has already identified as generally fitting his decision criteria. In this scenario, whether he chooses to take the decision as it stands or continue searching will depend on his expectations about finding another car that also fits his criteria and which might be better than those he has already seen. Thus, in this case, P3 refers to whether he will be able to: (i) acquire new information (e.g. that there is

such and such a car at such and such location), and (ii) make the necessary distinctions to frame that new information (i.e. that he will be able to relate the features of that vehicle to his decision criteria and correctly evaluate whether it is a viable candidate for his decision). Since our decision maker already knows approximately what he is looking for, and has already identified a number of other candidate vehicles for purchase, he is clearly not ignorant with respect to the future distinctions to which P3 refers; and indeed he may reasonably be confident about continuing to make those distinctions going forward.

More generally, the converse situation is also possible, where the agent is unsure of what distinctions might need to be made in the future, or if they will be able perform those distinctions, even when provided with the relevant information. An agent's state of belief in this regard will depend on whether he has any experience with making the relevant distinctions, and the agent's cognitive abilities (which influences how difficult those distinctions are for the agent to make). Expectations with respect to (i), above, will depend on the extent and duration of the agent's search to date (i.e. what and how much he already knows about the broader problem domain), the ease or difficulty of that search (as a guide to how easy or difficult it might be continue), and whether he has the time and resources to continue. Expectations with respect to (ii) above, will depend on the nature and reliability of the agent's knowledge of the broader problem domain, his experience with the specific framing process to date, and with the type of problem being addressed, and the nature and reliability of the agent's process of reasoning from this evidence to an expectation about the potential success of continuing the framing process. Table VIII.2 summarises these factors.

*Table VIII.2. Basis for belief in P3*

<p><b>Proposition 3:</b> That there are no other possible courses of action, which might be better than the best option already identified, and which could reasonably be identified within the constraints of the decision situation. Belief in P3 is based on an evaluation of the ease or difficulty of continuing to acquire information/knowledge, of making the relevant distinctions going forward, and of the potential for significant changes to the decision at hand.</p>	
<p><i>Belief in P3 is founded on the following evidence:</i></p> <ul style="list-style-type: none"> <li>Knowledge of the broader problem domain.</li> <li>Expectations about acquiring new information and knowledge about the problem domain.</li> <li>Expectations about making the necessary distinctions going forward.</li> <li>Expectations about the potential significance of those distinctions with respect to the decision at hand.</li> </ul>	<p><i>Degree of belief in P3 depends on:</i></p> <ul style="list-style-type: none"> <li>Reliability of knowledge about the present state of the world/system, and of the analytical framework/methodology used to make predictions in that system.</li> <li>Experience with the type of problem being addressed (making relevant distinctions), and with the framing process so far.</li> <li>Reliability of the agent's cognitive abilities.</li> </ul>

## General formulation of the satisficing decision

The preceding section elaborated the conditions for making the satisficing cut-off in relation to a particular set of circumstances. The assumed scale was, however, merely convenient for purposes of conceptualising the cut-off process; it was simply easier to think about a common decision task, and to describe the elements of the decision as aggregate wholes (i.e. objectives, preferences, options, outcomes, etc.). But while it is common, and therefore easier, to think about a decision in these terms, it is also, essentially, an arbitrary way of categorising decision elements. Objectives, preferences, options, and outcomes are merely labels for entities which have been identified, related, and grouped in a certain way. There is no reason why what I have called the “satisficing cut-off” should apply only once these aggregate categories have been established. Rather, the cut-off decision can be seen to also apply to every constitutive act of framing; that is, to every act of extending a frame to include a new entity or aspect thereof.

## General propositions and decision rules

Whenever we frame an entity, whenever we perceive the entity and distinguish it as somehow significant and relevant to some cognitive or practical endeavour, then if that framing is to be reasonable we must necessarily establish a degree of belief with respect to the following propositions (I use the shorthand notation, P, below to distinguish these general case propositions from those special case propositions, P, above):

- **Proposition 1 (P1) – Existence and nature of the entity:** That the entity exists and has certain aspects and characteristics as perceived.
- **Proposition 2 (P2) – Relevance and significance of the entity:** That the entity is significant and relevant with respect to what has already been framed.

At any given point in time, regardless of our state of belief with respect to P1 and P2, whether we seek to extend our frame depends on our state of belief with respect to the following proposition, if that extension is to be reasonable:

- **Proposition 3 (P3) – Value of extension:** That, going forward, new information and knowledge about the problem domain can be acquired, that new distinctions can be made, and that those distinctions will alter the decision at hand in some significant way.

For P1 and P2 there are two possible outcomes. Either the agent can justify a degree of



belief with respect to the truth or falsity of P1 and P2, with respect to any particular option; or the agent cannot justify belief either way. The agent may proceed in each case as follows (the numbering continues from the earlier list of decision rules):

- 5) If the agent cannot establish a degree of belief either way with respect to P1 or P2 then this is equivalent to a state of ignorance (i.e. insufficient knowledge) for the agent. The agent cannot formulate a reliable perception as to the existence and/or nature of the entity, and thus cannot judge whether the entity is significant and relevant, or not. In this case the agent may proceed as follows depending on his state of belief in P3:
  - a. Ignore the entity altogether (i.e. assume it does not exist) or bring the entity within the frame (i.e. in a hypothetical sense), but in either case settle the frame at that point. Whether the agent includes or excludes the entity will depend on how the agent deals with ignorance;
  - b. Either ignore the entity or include it within the frame, but in either case continue framing (i.e. continue the search for information and other entities to frame).
- 6) If the agent can establish a degree of belief with respect to the truth or falsity of P1 and P2 then:
  - a. If the agent believes either P1 or P2 to be false then, with respect to P1 no framing of the entity would be possible, and with respect to P2 the entity should be framed\*. In this case the agent may choose to do either of the following depending on his state of belief in P3:
    - i. Settle his frame; or
    - ii. Continue the framing process (i.e. continue the search for information and other entities to frame);

\* *refer to the clarificatory comments below for an explanation of why the entity should not be excluded from the frame.*

Or

- b. If the agent believes both P1 and P2 to be true then with respect to P1 framing of the entity becomes possible, and with respect to P2 the entity should be framed. In this case the agent may choose to do either of the following depending on his state of belief in P3:

- i. Settle his frame; or
- ii. Continue the framing process (i.e. continue the search for information and other entities to frame).

For P3 there are also two possible outcomes. Either the agent can justify a degree of belief with respect to the truth or falsity of P3; or the agent cannot justify belief either way.

7) If the agent cannot establish a degree of belief either way, then this constitutes a state of ignorance (i.e. insufficient knowledge), at least for that agent, in which case the agent may choose to satisfice as follows. Either:

- a. If time, resource, or other limits have already or will soon be encountered then it may nevertheless be reasonable to settle the frame as it currently stands; or
- b. If time, resource, or other limits do not impose an immediate need to settle the frame then it may not be unreasonable to continue the framing process for the time being; or
- c. If time, resource, or other limits do not impose an immediate need to settle the frame then it may not be unreasonable to settle the frame anyway, while reserving the option of restarting the framing process depending on what happens (i.e. proceeding on a trial and error basis).

8) If the agent can establish a degree of belief with respect to the truth or falsity of P3 then the agent may satisfice as follows. Either:

- a. If the agent believes P3 to be false then he should settle the frame as it stands;  
or
- b. If the agent believes P3 to be true then:
  - i. If time, resource, or other limits have already or will soon be encountered then it may nevertheless be reasonable to settle the frame as it stands; or
  - ii. If time, resource, or other limits do not impose an immediate need to settle the frame then the agent should continue the framing process for the time being.

## Qualifications and clarifications

The clarificatory comments made earlier with respect to belief in P1, P2, and P3 remain valid with respect to belief in P1, P2, and P3. These are summarised in the points below:

- Nothing is implied with respect to what constitutes justification in any specific sense. The objectivity of evidence for P1, P2, and P3 may be evaluated in a variety of ways; which methodologies and standards apply and how much justification is required depends on the nature of the evidence and who needs to be convinced.
- The evaluation of P3 by the agent is a continuous process, regardless of the agent's state of belief with respect to P1 and P2 for any given entity. The agent continuously evaluates whether to continue or cut-off the framing process.
- If the agent can establish a degree of belief with respect to the truth or falsity of P2 for any particular entity then, regardless of whether the agent believes P2 to be true or false, this implies that the agent has already evaluated the significance and relevance of the entity, accounting, presumably, for uncertainty where necessary. If this is the case then, regardless of whether the agent believes P2 to be true or false, the entity has already been framed; although the extent to which the agent is able to perform the evaluation, and their subsequent degree of belief, will of course depend on how much they know about the entity (in other words P2 depends on P1). In this sense, the act of forming a belief either way with respect to P2 constitutes the framing act. This means that the entity cannot simply be excluded from the frame if the agent comes to believe P2 to be false, since this would imply somehow returning the entity to the undistinguished space beyond the frame. Rather, to the extent that the entity has already been identified and evaluated it remains framed. Thus, the nature of the agent's belief with respect to P1 and P2 does not determine whether an entity becomes framed or not, but rather whether, having been framed, it should be categorised as relevant or irrelevant for the cognitive or practical endeavours at hand.
- The agent's state of belief with respect to P3 does not depend on any kind of assessment of things which might be called "potential" entities. If the agent were in a position to evaluate the relevance and significance of an entity then that entity would by definition already have been framed. Rather, the process of framing, of bringing an entity within the decision frame, involves drawing distinctions about that entity and its relationship to the problem or decision at hand. The important concept here is that what is or has been framed is distinguished from all else that is known by the agent, and from information in

general, by the judgements of relevance (or irrelevance) and significance (or insignificance). Proposition 3 relates to the agent's expectations about continuing to make such distinctions going forward, and about whether this might significantly alter the decision at hand (i.e. would those distinctions lead to a significantly different choice?)

### **Basis for belief in P1, P2, and P3**

Table VIII.3 lists the factors on which the agent's state and degree of belief with respect to P1, P2, and P3 necessarily depend.

*Table VIII.3. Basis for belief in P1, P2, and P3*

<b>Proposition 1 (P1):</b> That the entity exists and has certain aspects and characteristics as perceived. Belief in <u>P1</u> is based on an analysis of the nature and behaviour of the entity and its context.	
<i>Belief in <u>P1</u> is founded on the following evidence:</i> Knowledge of the entity including relationships with other entities and systems, and expectations about the future behaviour of the entity and its context.	<i>Degree of belief in <u>P1</u> depends on:</i> Reliability of knowledge about the entity and the world/system in which it is constituted, and of the analytical framework/methodology used to make predictions in and about that system.
<b>Proposition 2 (P2):</b> That the entity is significant and relevant with respect to what has already been framed. Belief in <u>P2</u> is based on an evaluation of the entity with respect to entities already framed, and against the criteria which define relevance and significance with respect to the frame.	
<i>Belief in <u>P2</u> is founded on the following evidence:</i> Knowledge of the entity, and other entities already framed. Knowledge of the criteria which define relevance and significance to the frame. A methodology for relating and comparing entities.	<i>Degree of belief in <u>P2</u> depends on:</i> Reliability of knowledge of entities Reliability of knowledge of the criteria which define relevance and significance to the frame. Reliability of the evaluative methodology
<b>Proposition 3 (P3):</b> That, going forward, new information and knowledge about the problem domain can be acquired, that new distinctions can be made and that those distinctions will alter the decision at hand in some significant way. Belief in <u>P3</u> is based on an evaluation of the ease or difficulty of continuing to acquire information/knowledge, of making the relevant distinctions going forward, and of the potential for significant changes to the decision at hand.	
<i>Belief in <u>P3</u> is founded on the following evidence:</i> Knowledge of the broader problem domain and expectations about acquiring new information and knowledge about the problem domain. Expectations about making the necessary distinctions going forward. Expectations about the potential significance of those distinctions with respect to the decision at hand.	<i>Degree of belief in <u>P3</u> depends on:</i> Reliability of knowledge about the present state of the world/system, and of the analytical framework/methodology used to make predictions in that system. Experience with the type of problem being addressed and the framing process so far. Reliability of the agent's cognitive abilities.

# References

- Aabo, T., J. R. S. Fraser, and B. J. Simkins. 2005. The Rise and Evolution of the Chief Risk Officer: Enterprise Risk Management at Hydro One. *Journal of Applied Corporate Finance* 17 (3):62-75.
- Abbate, J. 1994. The Internet Challenge: Conflict and Compromise in Computer Networking. In J. Summerton (ed), *Changing Large Technical Systems*. Boulder, Co.: Westview Press, 193-210.
- Abbott, A. 1988. *The System of Professions: An Essay on the Division of Expert Labour*. Chicago: The University of Chicago Press.
- Acharyya, M. 2008. An Empirical Study on Enterprise Risk Management in Insurance. In D. L. Olson and D. Wu (eds), *New Frontiers in Enterprise Risk Management*. Berlin, Heidelberg: Springer, 39-55.
- Acharyya, M., and E. J. Johnson. 2006. Investigating the development of ERM in the insurance industry: An empirical study of four major European insurers. *The Geneva Papers on Risk and Insurance: Issues and Practice* Special Issue, July:55-80.
- Adler, N., A. B. Shani, and A. Styhre (eds). 2004. *Collaborative Research in Organizations: Foundations for Learning, Change, and Theoretical Development*. Thousand Oaks, Ca.: Sage Publications.
- AIRMIC, ALARM, and IRM. 2002. *A Risk Management Standard*. London: Association of Insurance and Risk Managers (AIRMIC), The Association of Local Authority Risk Managers (ALARM), and the Institute of Risk Management (IRM). Available online at <http://www.alarm-uk.org/publications.aspx>.
- Alferoff, C., and D. Knights. 2009. Making and Mending your Nets: Managing Relevance, Participation and Uncertainty in Academic-Practitioner Knowledge Networks. *British Journal of Management* 20 (2009):125-142.
- Althaus, C. E. 2005. A Disciplinary Perspective on the Epistemological Status of Risk. *Risk Analysis* 25 (3):567-588.
- Alvesson, M., and D. Kärreman. 2001. Odd couple: Making sense of the curious concept of knowledge management. *Journal of Management Studies* 38 (7):995-1018.
- Andrews, K. R. 1980. *The concept of corporate strategy*, Revised edition. Homewood, Il.: Richard D. Irwin Inc.
- Ang, A. H.-S., and D. De Leon. 2005. Modelling and analysis of uncertainties for risk-informed decisions in infrastructures engineering. *Structure and Infrastructure Engineering* 1 (1):19-31.

- AON Corporation. 2009. Global Risk Management Survey '09. Published online by AON Corporation, [www.aon.com](http://www.aon.com).
- Aram, J. D., and P. F. Salipante. 2003. Bridging Scholarship in Management: Epistemological Reflections. *British Journal of Management* 14 (2003):189-205.
- Argyris, C., and D. A. Schön. 1974. *Theory in Practice*. San Francisco, Ca.: Jossey-Bass.
- . 1978. *Organizational Learning: A Theory of Action Perspective*. Reading, MA.: Addison-Wesley Publishing.
- Armstrong, J. 2008. *Design Matters: The Organisation and Principles of Engineering Design*. London: Springer-Verlag.
- Atkinson, W. 2007. A View from the Top: The Growing Role and Influence of the Chief Risk Officer. *Risk Management Magazine* 2007 (September):24-30.
- Auckland Water Group. 2007. *Auckland Water Industry Annual Performance Review 2006/07*. Auckland: Auckland Water Group.
- Auyang, S. Y. 2004. *Engineering: an endless frontier*. Cambridge, Ma.: Harvard University Press.
- Aven, T. 2004. On How to Approach Risk and Uncertainty to Support Decision-making. *Risk Management: An International Journal* 6 (4):27-39.
- Ayyub, B. M. 2003. *Risk Analysis in Engineering and Economics*. Boca Raton, Fl.: Chapman & Hall/CRC.
- Baert, P. 2003. Pragmatism, Realism and Hermeneutics. *Foundations of Science* 8:89-106.
- . 2005. Towards a Pragmatist-Inspired Philosophy of Social Science. *Acta Sociologica* 48 (3):191-203.
- Baime, A. J. 2009. A War of Speed. *Playboy Magazine* 56 (5):34-36 & 96-100.
- Barton, T. L., W. G. Shenkir, and P. L. Walker. 2002. *Making Enterprise Risk Management Pay Off*. Upper Saddle River, NJ.: Financial Times/Prentice Hall.
- Basel Committee on Banking Supervision. 2004. Working Paper No. 13: Bank Failures in Mature Economies. Basel, Switzerland: Bank for International Settlements (available online at <https://www.bis.org>).
- Beasley, M., D. Pagach, and R. Warr. 2008. Information conveyed in hiring announcements of senior executives overseeing enterprise-wide risk management processes. *Journal of Accounting, Auditing & Finance* 23 (3):311-332.
- Beasley, M. S., B. C. Branson, and B. V. Hancock. 2009. ERM: Opportunities for Improvement. *Journal of Accountancy* 205 (4):44-51.
- Beasley, M. S., R. Clune, and D. R. Hermanson. 2005a. Enterprise risk management: An

- empirical analysis of factors associated with the extent of implementation. *Journal of Accounting and Public Policy* 24 (2005):521-531.
- . 2005b. ERM: A Status Report. *Internal Auditor* 62 (1):67-72.
- Beaumont, P. B., J. W. Leopold, and J. R. Coyle. 1982. The safety officer: an emerging management role? *Personnel Review* 11 (2):35-38.
- Beck, M. 2004. Obstacles to the Evolution of Risk Management as a Discipline: Some Tentative Thoughts. *Risk Management: An International Journal* 6 (3):13-21.
- Becker, F. S. 2006. Globalization, curricula reform and the consequences for engineers working in an international company. *European Journal of Engineering Education* 31 (3):261-272.
- Beer, S. 1972. *Brain of the Firm*. London: Penguin.
- . 1985. *Diagnosing the System for Organizations*. London and New York: Wiley.
- Bell, D. 1999. The axial age of technology foreword: 1999. In *The Coming of Post Industrial Society*. New York: Basic Books, Special Anniversary Edition, pp ix-lxxxv,
- Bernstein, P. L. 1996. *Against the gods: the remarkable story of risk*. New York: John Wiley & Sons.
- Bertho-Lavenir, C. 1988. The telephone in France 1879 to 1979: National characteristics and international influences. In R. Mayntz and T. P. Hughes (eds), *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press, 155-178.
- Bijker, W. E., T. P. Hughes, and T. J. Pinch (eds). 1987. *The Social Construction of Technological Systems: new directions in the sociology and history of technology*. Cambridge, MA.: MIT Press.
- Blumer, H. 1969. *Symbolic interactionism: Perspective and method*. Berkeley, Los Angeles: University of California Press.
- Bonazzi, G. 1998. Between Shock Absorption and Continuous Improvement: Supervisors and Technicians in a Fiat "Integrated Factory". *Work Employment and Society* 12 (2):219-243.
- Bosch, R. 2007. Pragmatism and the practical relevance of truth. *Foundations of Science* 12 (2007):189-201.
- Bosco, F. J. 2006. Actor-Network Theory, Networks, and Relational Approaches in Human Geography. In S. Aitken and G. Valentine (eds), *Approaches to Human Geography*. London: Sage Publications, 136-146.
- Bowden, A. R., M. R. Lane, and J. H. Martin. 2001. *Triple Bottom Line Risk Management: Enhancing Profit, Environmental Performance, and Community Benefits*. New York: John Wiley & Sons, Inc.

- Bowker, G. C., and S. L. Star. 1999. *Sorting Things Out: Classification and its Consequences*. Cambridge, MA: Massachusetts Institute of Technology.
- Boyle, C. 2002. Sustainability Assessment: Assessing the Sustainability of Business Activities. Corporate Social Responsibility & Governance for Sustainability, 10th International Conference of the Greening of Industry Network, June 23 - 26, Chalmers University of Technology, Göteborg, Sweden. The Greening of Industry Network, available at <http://gin.confex.com/gin/archives/2002/papers/010249Boyle.pdf>.
- . 2004a. Achieving Sustainability. 1st International Conference on Sustainability Engineering and Science, 6 - 9 July, Auckland, New Zealand. The New Zealand Society for Sustainability Engineering and Science, available at <http://www.nzsses.auckland.ac.nz/conference/2004/index.htm>.
- . 2004b. Achieving Sustainability. International Centre for Sustainability Engineering and Research, University of Auckland: Unpublished.
- Brancato, C. K., M. Tonello, E. Hexter, and K. R. Newman. 2006. *The Role of U.S. Corporate Boards in Enterprise Risk Management*. New York: The Conference Board.
- Brim, O. G. J., D. C. Glass, D. E. Lavin, and N. Goodman. 1962. *Personality and decision processes: Studies in the Social Psychology of Thinking*. Stanford: Stanford University Press.
- Bronet, F., R. Eglash, G. Gabriele, D. Hess, and L. Kagan. 2003. Product Design and Innovation: Evolution of an Interdisciplinary Design Curriculum. *International Journal of Engineering Education* 19 (1):183-191.
- Brown, J. S., and P. Duguid. 1991. Organizational Learning and Communities-of-Practice: Toward a Unified View of Working, Learning, and Innovation. *Organization Science* 2 (1):40-57.
- . 2002. *The Social Life of Information*. Boston, MA.: Harvard Business School Press.
- Brown, M. 1997. A CFO's view: "two-feet-on-the-ground" risk management. *Risk Management* 1997 (September ):21-27.
- Buehler, K., A. Freemant, and R. Hulme. 2008. The New Arsenal of Risk Management. *Harvard Business Review* 86 (9):92-100.
- Buenaño, G. 1999. *The Becoming of Problems in Design: Knowledge in Action to Frame Wicked Problems*, PhD Thesis, Graduate Division of the University of California, Berkeley.
- Burlando, A. J. 1990. The 1990s: The Decade of Risk Management. *Risk Management* 37 (3):50-56.
- Burrell, G., and G. Morgan. 1979. *Sociological Paradigms and Organizational Analysis: Elements of the Sociology of Corporate Life*. London: Heinemann.
- Cabantous, L., and J.-P. Gond. 2006. (Re)Constructing Rationality in Organizations: Decision-



- Making as a 'Performative Praxis'. Unpublished research paper available online the Centre for Decision Research, University of Leeds, UK, at <http://lubswww2.leeds.ac.uk/cdr/index.php?id=154>: [accessed 7 February 2007].
- Cabantous, L., J.-P. Gond, and M. Johnson-Cramer. 2008. The social construction of rationality in organizational decision making. In G. P. Hodgkinson and W. H. Starbuck (eds), *The Oxford Handbook of Organizational Decision Making*. Oxford: Oxford University Press, 399-417.
- Callon, M. 1998a. An essay on framing and overflowing: economic externalities revisited by sociology. In Oxford, UK: Blackwell Publishers/*The Sociological Review*, 244-269.
- . 1998b. Introduction: The embeddedness of economic markets in economics. In M. Callon (ed), *The Laws of the Markets*. Oxford, UK: Blackwell Publishers/*The Sociological Review*, 1 - 57.
- . 1999. Actor-network theory - the market test. In J. Law and J. Hassard (eds), *Actor Network Theory and After*. Oxford: Blackwell Publishers/*The Sociological Review*, 181-195.
- (ed). 1998c. *The Laws of the Markets*. Oxford, UK: Blackwell Publishers/*The Sociological Review*.
- Callon, M., and B. Latour. 1981. Unscrewing the big Leviathan: how actors macro-structure reality and how sociologists help them to do so. In K. Knorr-Cetina and A. V. Cicourel (eds), *Advances in social theory and methodology: Toward an integration of micro- and macro-sociologies*. Boston, MA.: Routledge & Kegan Paul, 277-303.
- Callon, M., and J. Law. 1997. After the Individual in Society: Lessons on Collectivity from Science, Technology and Society. *Canadian Journal of Sociology* 22 (2):165-182.
- Callon, M., and F. Muniesa. 2005. Economic Markets as Calculative Collective Devices. *Organization Studies* 26 (8):1229-1250.
- Capability Maturity Model. 2008. In Wikipedia, The Free Encyclopedia. Retrieved 17 December 2008, from [http://en.wikipedia.org/w/index.php?title=Capability\\_Maturity\\_Model&oldid=258202966](http://en.wikipedia.org/w/index.php?title=Capability_Maturity_Model&oldid=258202966).
- Cardwell, D. 1994. *The Fontana History of Technology*. London: Fontana Press.
- Caron, F. 1988. The evolution of the technical system of railways in France from 1832 to 1937. In R. Mayntz and T. P. Hughes (eds), *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press, 69-104.
- Carroll, A. B., and A. K. Buchholtz. 2000. *Business & society: ethics and stakeholder management*. London: International Thomson.
- Castree, N. 2006. Commentary. From neoliberalism to neoliberalisation: consolations, confusions, and necessary illusions. *Environment and Planning A* 38 (2006):1-6.

- Chaiken, S., and Y. Trope (eds). 1999. *Dual-Process Theories in Social Psychology*. New York: Guilford Press.
- Champion, D., R. S. Kaplan, A. Mikes, R. Simons, P. Tufano, and M. Hofman. 2009. Managing risk in the new world. *Harvard Business Review* 87 (1):68-75.
- Champlin, J. B. 2004. Beyond the CMM: Why Implementing the SEI's Capability Maturity Model is Insufficient to Deliver Quality Information Systems in Real-World Corporate IT Organizations. Presented to the Information Systems Audit and Control Association (ISACA), January 2004, Chicago. Available online at <http://faculty.roosevelt.edu/Champlin/>.
- Chermack, T. J. 2005. Studying scenario planning: Theory, research suggestions, and hypotheses. *Technological Forecasting & Social Change* 72 (2005):59-73.
- Chia, R. 1994. The Concept of Decision: A Deconstructive Analysis. *Journal of Management Studies* 31 (6):781-806.
- Choo, C. W. 2006. *The Knowing Organization*. 2nd ed. New York: Oxford University Press.
- Ciccarelli, M. C. 2001. CROs: The Rising Tide. *Risk & Insurance* 12 (4):15-18.
- . 2002. New Skills for New Challenges. *Risk & Insurance* 13 (4):16-20.
- . 2003. CROs: is the glitz going...going...gone? *Risk & Insurance* 14 (6):14-16.
- Clement, J. 2007a. Establishing the Risk Context (Draft). Internal paper (November 2007).
- . 2007b. Rationale for consequence framework (Internal paper August 2007).
- . 2007c. Staff Survey WSL Risk Management, August - September 2007, Summary Report. Internal paper (October 2007).
- . 2007d. WSL Risk Management - The Concept. Internal paper (October 2007).
- . 2007e. WSL Risk Management Framework Development Plan 2007 - 2010. Internal paper (March 2007).
- . 2007f. WSL Risk Management Vision. Internal paper (May 2007).
- . 2008a. Establishing the Risk Context (Version 2.0). Internal paper (February 2008).
- . 2008b. In search of the Holy Grail - practical challenges in the pursuit of ERM. 6th Annual Enterprise Risk Management Conference, 28 - 29 April 2008, Auckland, New Zealand. Conferenz Ltd.
- . 2008c. Risk Identification: Foundation Framework. Internal paper (May 2008). Watercare Services Ltd.
- . 2008d. Risk Identification: Foundation Framework. Internal paper (May 2008).
- . 2008e. Strategic Risk Assessment. Internal paper (March 2008).

- . 2008f. WSL Risk Management Business Operations Risks: Consequences Framework (WIP 20 February 2008).
- Coffin, B. 2009a. The Way Forward: Rethinking Enterprise Risk Management. *Risk Management* 56 (3):36-40.
- (ed). 2009b. The 2008 Financial Crisis: A Wake-up Call for Enterprise Risk Management. RIMS Executive Report. Risk and Insurance Management Society. Published online at <http://community.rims.org/RIMS/RIMS/Community/Resources/ViewDocument/Default.aspx?DocumentKey=6742a2ba-1b1d-446e-85c3-b0c8880f42fb>.
- Colquitt, L. L., R. E. Hoyt, and R. B. Lee. 1999. Integrated Risk Management and the Role of the Risk Manager. *Risk Management and Insurance Review* 2 (3):43-61.
- Committee on the Financial Aspects of Corporate Governance. 1992. Report of the Committee on the Financial Aspects of Corporate Governance,. London: Gee and Co. Ltd, available online at the European Corporate Governance Institute, <http://www.ecgi.org/codes/documents/cadbury.pdf>.
- Conference Board of Canada. 2001. A Composite Sketch of a Chief Risk Officer. Ottawa: Conference Board of Canada.
- Conley, J. 2000. Leaders of the evolution. *Risk Management* 47 (12):43-50.
- Constant, E. W. 1999. Reliable Knowledge and Unreliable Stuff: On the Practical Role of Rational Beliefs. *Technology and Culture* 40 (April):324-357.
- Cook, S. D. N., and J. S. Brown. 1999. Bridging Epistemologies: The Generative Dance Between Organizational Knowledge and Organizational Knowing. *Organization Science* 10 (4):381-400.
- Cooper, D. F., G. Raymond, J. Hodges, and L. Martin. 1998. Draft Risk Management Framework. Auckland: Report prepared for Watercare Services Ltd by Broadleaf Capital International and Woodward-Clyde New Zealand.
- Cooper, R. 1986. Organization/Disorganization. *Social Science Information* 25 (2):299-335.
- Corbett, R. B. 2004. A View of the Future of Risk Management. *Risk Management: An International Journal* 6 (3):51-56.
- COSO. 1991. Internal Control - Integrated Framework. Committee of Sponsoring Organizations of the Treadway Commission (COSO). Executive Summary available online at <http://www.coso.org>.
- . 2004. Enterprise Risk Management - Integrated Framework. Committee of Sponsoring Organizations of the Treadway Commission (COSO). Executive Summary available online at <http://www.coso.org>.
- Costanza, R. 2006. Limits to growth: The 30-year update. *Ecological Economics* 59 (3):397-399.

- Costanza, R., and H. E. Daly. 1987. Toward an ecological economics. *Ecological Modelling* 38 (1-2):1-7.
- . 1992. Natural capital and sustainable development. *Conservation Biology* 6 (1):37-46.
- Covello, V. T., and J. Mumpower. 1985. Risk Analysis and Risk Management: An Historical Perspective. *Risk Analysis* 5 (2):103-120.
- Creswell, J. W. 1998. *Qualitative inquiry and research design: choosing among five traditions*. Thousand Oaks, CA.: Sage.
- Crosby, P. B. 1979. *Quality is free: the art of making quality certain*. New York: McGraw-Hill.
- Culp, C. L. 2002. The revolution in corporate risk management: a decade of innovations in process and products. *Journal of Applied Corporate Finance* 14 (4):8-26.
- Cummings, G. 2008. Convergence: An Integrated Risk and Governance Model. Panel discussion at the 6th Annual Enterprise Risk Management Conference, 28 - 29 April 2008, Auckland, New Zealand. Conferenz Ltd.
- Czarniawska, B. 2008. Organizing: how to study it and how to write about it. *Qualitative Research in Organizations and Management: An International Journal* 3 (1):4-20.
- D'Elia, A. 2002. Pursuing the Vision Thing. *Risk & Insurance* 12 (10):42.
- Dai, Y.-S., M. Xie, Q. Long, and S.-H. Ng. 2007. Uncertainty Analysis in Software Reliability Modeling by Bayesian Approach with Maximum-Entropy Principle. *IEEE Transactions on Software Engineering* 33 (11):781-795.
- Dane, E., and M. G. Pratt. 2007. Exploring intuition and its role in managerial decision making. *Academy of Management Review* 32 (1):33-54.
- Davies, A. 1996. Innovation in Large Technical Systems: The Case of Telecommunications. *Industrial and Corporate Change* 5 (4):1143-1180.
- de Bruijn, T., and V. Norberg-Bohm. 2005. Introduction: Toward a New Paradigm for the Transition to a Sustainable Society? In T. de Bruijn and V. Norberg-Bohm (eds), *Industrial Transformation. Environmental Policy Innovation in the United States and Europe*. Cambridge, Ma.: MIT Press, 1-36.
- de Camp, L. S. 1977. *Ancient Engineers*. London: Tandem.
- de Graaf, E., and W. Ravesteijn. 2001. Training complete engineers: global enterprise and engineering education. *European Journal of Engineering Education* 26 (4):419-427.
- de la Rosa, S. 2007. Taking a closer LOOK at the role Chief Risk Officer. *Accountancy SA* 2007 (March):34-35.
- DeLoach, J. W. 2000. *Enterprise-Wide Risk Management*. London: Financial Times-Prentice Hall.

- Deloitte and Touche LLP. 2004. 2004 Global Risk Management Survey. Published online at <http://www.deloitte.com>: Deloitte and Touche LLP.
- . 2007. Global Risk Management Survey: Fifth Edition - Accelerating Risk Management Practices. Published online at <http://www.deloitte.com>: Deloitte and Touche LLP.
- Desender, K. A. 2007. The influence of board composition on enterprise risk management implementation. Working paper. Available at [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1025982](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1025982).
- Dewey, J. 1910. How we think. Boston: D.C. Heath & Co.
- Diamond, J. 2005. Collapse - How Societies Choose to Fail or Succeed. New York: Viking Penguin.
- Dickinson, G. 2001. Enterprise Risk Management: Its Origins and Conceptual Foundation. *The Geneva Papers on Risk and Insurance* 26 (3):360-366.
- Dixon, G. 2009. Taking Risk Management Back Into The Boardroom. *The Corporate Board* 30 (179):17-21.
- Donnelly, R., and C. Boyle. 2004. Long-term planning for the sustainability of water service provision and infrastructure for the Auckland Region. Technology New Zealand TIF Application, July 2004.
- Douglas, M. 1985. Risk Acceptability According to the Social Sciences. New York: Russel Sage Foundation.
- Dreyer, S. J., and D. Ingram. 2007. Request For Comment: Enterprise Risk Management Analysis For Credit Ratings Of Nonfinancial Companies. *RatingsDirect* 2007 (November 15):1-11, published by Standard & Poors, available online at [http://www.rims.org/Content/NavigationMenu/ERM/Papers\\_and\\_Studies/ERM\\_ratingspracticesSandP.pdf](http://www.rims.org/Content/NavigationMenu/ERM/Papers_and_Studies/ERM_ratingspracticesSandP.pdf).
- Drucker, P. 1969. The age of discontinuity: guidelines to our changing society. London: Heinemann.
- Dubin, R. 1978. Theory Building (revised ed.). New York: Free Press/MacMillan.
- Duncan, R. B. 1979. What is the Right Organization Structure? Decision Tree Analysis Provides the Answer. *Organizational Dynamics* 7 (3):59-80.
- Dym, C. L. 2008. Educating Engineers for a Flat World. *International Journal of Engineering Education* 24 (2):214-220.
- Economist Intelligence Unit. 2005. The evolving role of the CRO. Available online at <http://www.eiu.com>: Economist Intelligence Unit.
- . 2009. Managing risk in perilous times: Practical steps to accelerate recovery. Available online at <http://www.eiu.com> (March 2009): Economist Intelligence Unit.

- Eick, C. L. M. 2003. Factors that promote effective risk management at universities classified by the Carnegie system as Doctoral/Research Universities - Extensive, Diss. for Doctor of Education, the Graduate Faculty, Auburn University, Auburn, Alabama.
- Einstein, A. 1954. Ideas and Opinions. New York: Crown.
- Ekins, P. 2003. Identifying critical natural capital: Conclusions about critical natural capital. *Ecological Economics* 44 (2-3):277-292.
- Ekins, P., and S. Simon. 2003. An illustrative application of the CRITINC framework to the UK. *Ecological Economics* 44 (2-3):255-275.
- Ekins, P., S. Simon, L. Deutsch, C. Folke, and R. De Groot. 2003. A framework for the practical application of the concepts of critical natural capital and strong sustainability. *Ecological Economics* 44 (2-3):165-185.
- Emblemsvåg, J., and B. Bras. 2000. Process thinking - a new paradigm for science and engineering. *Futures* 32 (2000):635-654.
- epistemology. 1999. The Cambridge Dictionary of Philosophy. Retrieved August 26, 2010, from <http://www.credoreference.com.ezproxy.auckland.ac.nz/entry/cupdphil/epistemology>.
- Ernst & Young. 2006a. Board members on risk. Leveraging frameworks for the future. Published online at <http://www.ey.com>: Ernst & Young.
- . 2006b. Companies on risk. The benefits of alignment. Available online at <http://www.ey.com>: Ernst & Young.
- . 2007. Managing risk across the enterprise: Building a comprehensive approach to risk. Available online at <http://www.ey.com>: Ernst & Young.
- . 2008. Moving to the next level - a progress report on insurance risk leadership. 2008 Insurance Risk Leadership Survey. Published online at <http://www.ey.com>: Ernst & Young.
- . 2009. The future of risk: Protecting and enabling performance. Published online at <http://www.ey.com>: Ernst & Young.
- Espejo, R., and R. Harnden. 1989. The Viable Systems Model - Interpretations and Applications of Stafford Beer's VSM. London and New York: Wiley.
- Evans, J. S. B. T. 2007. Hypothetical Thinking: Dual Processes in Reasoning and Judgement. Hove, UK.: Psychology Press.
- . 2008. Dual-Processing Accountings of Reasoning, Judgement, and Social Cognition. *Annual Review of Psychology* 59:255-278.
- Falkenstein, E. 2001. The Risk Manager of the Future: Scientist or Poet? *The RMA Journal* 2001 (February):18-22.

- Feilzer, M. Y. 2010. Doing Mixed Methods Research Pragmatically: Implications for the Rediscovery of Pragmatism as a Research Paradigm. *Journal of Mixed Methods Research* 4 (1):6-16.
- Fendt, J., R. Kaminska-Labbé, and W. M. Sachs. 2007. Concepts of Truth and Revelance in Management Research: a Pragmatic Consideration. The Third Organization Studies Summer Workshop on "Organization Studies as Applied Science: The Generation and Use of Academic Knowledge about Organizations". Crete, Greece, June 7-9.
- Fenton, N., D. Deacon, A. Bryman, and P. Birmingham. 1997. Sod Off and Find Us a Boffin: Journalists and the Social Science Conference. *The Sociological Review* 45 (1):1-23.
- FIDIC. 1990. Consulting engineers and the environment. FIDIC Policy Statement, June 1990. International Federation of Consulting Engineers, available online at <http://www1.fidic.org/about/statement04.asp>.
- . 2002. Sustainable development in the consulting engineering industry: A unique capacity to address the priorities. Lausanne, Switzerland: International Federation of Consulting Engineers.
- Fischer, H. R. 2001. Abductive reasoning as a way of worldmaking. *Foundations of Science* 6 (2001):361-383.
- Fischhoff, B., S. R. Watson, and C. Hope. 1984. Defining Risk. *Policy Sciences* 17 (1984):123-139.
- Fogg, E. 2006. CROs - A Fad or the Future? *Risk & Insurance* 17 (November):12.
- Frame, J. D. 2003. *Managing Risk in Organizations: A Guide for Managers*. San Francisco, CA.: Jossey-Bass.
- Fraser, J., and B. J. Simkins. 2007. Ten Common Misconceptions About Enterprise Risk Management. *Journal of Applied Corporate Finance* 19 (4):75-81.
- (eds). 2010. *Enterprise Risk Management: Today's Leading Research and Best Practices for Tomorrow's Executives.*, The Robert W. Kolb Series in Finance, Hoboken, New Jersey: John wiley & Sons, Inc.
- Fraser, J. R. S., K. Schoening-Thiessen, and B. J. Simkins. 2008. Who Reads What Most Often? A Survey of Enterprise Risk Management Literature Read by Risk Executives. *Journal of Applied Finance* 18 (1):73-91.
- Frost, R. B. 1992. A Converging Model of the Design Process: Analysis and Creativity, the Ingredients of Synthesis. *Journal of Engineering Design* 3 (2):117-126.
- Funtowicz, S. O., and J. R. Ravetz. 1993. Science for the post-normal age. *Futures* 25 (7):739-755.
- Galambos, L. 1988. Looking for the boundaries of technological determinism: A brief history of the U.S. telephone system. In R. Mayntz and T. P. Hughes (eds), *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press, 135-154.

- Gardner, H. 1983. *Frames of mind: the theory of multiple intelligences*. New York: Basic Books.
- . 1999. *Intelligence Reframed: Multiple Intelligences for the 21st Century*. New York: Basic Books.
- Gates, S., J. L. Nicolas, and P. L. Walker. 2009. Perceived value of enterprise risk management. University of Virginia Working Paper.
- Geels, F. W. 2007. Transformations in Large Technical Systems: A Multilevel Analysis of the Dutch Highway System (1950-2000). *Science, Technology, & Human Values* 32 (2):123-149.
- Ghoshal, S. 2005. Bad management theories are destroying good management practices. *Academy of Management Learning & Education* 4 (1):75-91.
- Gilovich, T., D. Griffith, and D. Kahneman (eds). 2002. *Heuristics and Biases: The Psychology of Intuitive Judgement*. Cambridge: Cambridge University Press.
- Glaser, B. G., and A. L. Strauss. 1967 [2006]. *The discovery of grounded theory: strategies for qualitative research*. New Brunswick: Aldine Transaction.
- Gökalp, I. 1992. On the Analysis of Large Technical Systems. *Science, Technology, & Human Values* 17 (1):57-78.
- Gordon-Clark, J. 2007. Risky Business - The Evolving Role of the CRO. Available online at <http://www.heidrick.com/IC/APLCS.htm>: Heidrick and Struggles.
- Gordon, L. A., M. P. Loeb, and C.-Y. Tseng. 2009. Enterprise risk management and firm performance: A contingency perspective. *Journal of Accounting and Public Policy* 28 (4):301-327.
- Graziano, C., and R. Aggarwal. 2005. C-suites Gaining Another Member. *Financial Executive* September 2005:42-44.
- GRI. 2002. *Sustainability Reporting Guidelines*. Boston, MA.: Global Reporting Initiative (<http://www.globalreporting.org>).
- Gunderson, L., and C. S. Holling. 2001. *Panarchy: understanding transformations in human and natural systems*. Washington D.C.: Island Press.
- Hadorn, G. H., S. Biber-Klemm, W. Grossenbacher-Mansuy, H. Hoffman-Riem, D. Joye, C. Pohl, U. Wiesmann, and E. Zemp. 2008. The Emergence of Transdisciplinarity as a Form of Research. In G. H. Hadorn, H. Hoffman-Riem, S. Biber-Klemm, W. Grossenbacher-Mansuy, D. Joye, C. Pohl, U. Wiesmann and E. Zemp (eds), *Handbook of Transdisciplinary Research*. New York: Springer, 19-39 (Chapter 2).
- Hall, J., R. Dawson, L. Manning, M. Walkden, M. Dickson, and P. Sayers. 2006. Managing changing risks to infrastructure. *Proceedings of the Institute of Civil Engineers, Civil Engineering* 159 (November):21-27.



- Hansson, S. O. 1994. Decision Theory: A Brief Introduction (revised ed. 2005). Available online from the Royal Institute of Technology, Stockholm (<http://www.infra.kth.se/~soh/decisiontheory.pdf>). Downloaded 15 September 2010.
- Harrington, S. E., G. Niehaus, and K. J. Risko. 2002. Enterprise risk management: the case of United Grain Growers. *Journal of Applied Corporate Finance* 14 (4):71-81.
- Harris, J. 2000. Basic Principles of Sustainable Development, Working Paper 00-04. Global Development and Environment Institute, Tufts University, MA. Available online at <http://ase.tufts.edu/gdae>.
- Hatchuel, A. 2000. Intervention Research and the Production of Knowledge. In M. Cerf, D. Gibbon, B. Hubert, R. Ison, J. Jiggins, M. Paine, J. Proost and N. Röling (eds), *Cow up a tree: knowing and learning for change in agriculture*. Paris: INRA, 55-68.
- . 2001. The Two Pillars of New Management Research. *British Journal of Management* 12 (Special Issue):S33-S39.
- Hay, P. 2002. *Main currents in western environmental thought*. Sydney, Australia: University of New South Wales Press.
- HEFCE. 1998. Industry-Academic Links in the UK (98/70). Available online at [http://www.hefce.ac.uk/pubs/hefce/1998/98\\_70.htm](http://www.hefce.ac.uk/pubs/hefce/1998/98_70.htm): Higher Education Funding Council for England.
- Heinze, G. W., and H. H. Kill. 1988. The development of the German railroad system. In R. Mayntz and T. P. Hughes (eds), *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press, 105-134.
- Helfat, C. E., S. Finkelstein, W. Mitchell, M. A. Peteraf, H. Singh, D. J. Teece, and S. G. Winter. 2007. *Dynamic Capabilities: Understanding Strategic Change in Organizations*. Malden, MA: Blackwell.
- Herbst, P. G. 1976. Alternatives to Hierarchies. *International series on the quality of working life*, 1, H. van Beinum (ed). Leiden: Martinus Nijhoff Social Sciences Division.
- Hilgartner, S. 1992. The Social Construction of Risk Objects: Or, How to Pry Open Networks of Risk. In J. F. Short and L. Clarke (eds), *Organizations, Uncertainties, and Risk*. Boulder, CO.: Westview Press Inc., 39-56.
- Hirsh, S. G. 1999. Children's relevance criteria and information seeking on electronic resources. *Journal of the American Society for Information Science* 50 (14):1265-83.
- Hodgkinson, G. P., E. Sadler-Smith, L. A. Burke, G. Claxton, and P. R. Sparrow. 2009. Intuition in Organizations: Implications for Strategic Management. *Long Range Planning* 42 (2009):277-297.
- Hodgkinson, G. P., and W. H. Starbuck. 2008. Organizational decision making: mapping terrains on different planets. In G. P. Hodgkinson and W. H. Starbuck (eds), *The Oxford Handbook of Organizational Decision Making*. Oxford: Oxford University Press, 1-29.

- Hoffman, A. 2004. Reconsidering the role of the practical theorist: On (re)connecting theory to practice in organization theory. *Strategic Organization* 2 (2):213-222.
- Hookway, C. 2010. Pragmatism. The Stanford Encyclopedia of Philosophy (Spring 2010 Edition), ed. E. N. Zalta: <http://plato.stanford.edu/archives/spr2010/entries/pragmatism/>.
- Hopkin, P. 2002. *Holistic Risk Management in Practice*. London: Witherby.
- Hora, S. C. 1996. Aleatory and epistemic uncertainty in probability elicitation with an example from hazardous waste management. *Reliability Engineering and System Safety* 54 (1996):217-223.
- Hoyt, R. E., and A. P. Liebenberg. 2009. The Value of Enterprise Risk Management: Working Paper. Available at [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1440947](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1440947).
- Huff, A. S., and J. O. Huff. 2001. Re-Focussing the Business School Agenda. *British Journal of Management* 12 (Special Issue):S49-S54.
- Hughes, T. P. 1983. *Networks of power: electrification in Western society 1880-1930*. Baltimore and London: Johns Hopkins University Press.
- . 1987. The Evolutions of Large Technical Systems. In T. J. Pinch, W. E. Bijker and T. P. Hughes (eds), *The Social Construction of Technological Systems: new directions in the sociology and history of technology*. Cambridge, Ma.: MIT Press, 51-82.
- Hurst, K. 1999. *Engineering Design Principles*. London: Arnold.
- Hutchins, E. 1995. *Cognition in the Wild*. Cambridge, Ma.: MIT Press, Cambridge University.
- Hutter, B., and M. Power. 2005a. Organizational encounters with risk: an introduction. In B. Hutter and M. Power (eds), *Organizational Encounters with Risk*. Cambridge: Cambridge University Press, 1-32.
- (eds). 2005b. *Organizational Encounters with Risk*. Cambridge: Cambridge University Press.
- IACCM. 2003. *Organisational Maturity in Business Risk Management: The IACCM Business Risk Management Maturity Model (BRM3)*. IACCM Business Risk Management Working Group, available online at <http://www.risk-doctor.com/pdf-files/brm1202.pdf>.
- IBNET. 2004. *The IBNET Toolkit. The International Benchmarking Network for Water and Sanitation Utilities* (<http://www.ib-net.org>).
- Ignatius, A. 2009. The Trouble with Risk. *Harvard Business Review* 87 (1):16.
- Installé, M. 1996. How to Educate Future Engineers Towards a Better Understanding of the Relationships Between Technology, Society and the Environment? *European Journal of Engineering Education* 21 (4):341-345.

- IPENZ. 2007. Engineers and Ethical Obligations. Practice Note 08 (February 2007). Institute of Professional Engineers New Zealand. Available online at: [http://www.ipenz.org.nz/ipenz/forms/pdfs/PN08\\_Ethical\\_Obligations.pdf](http://www.ipenz.org.nz/ipenz/forms/pdfs/PN08_Ethical_Obligations.pdf).
- . 2009a. Graduate Competency Profiles. Version 12, November 2009. Institute of Professional Engineers New Zealand. Available online at: [http://www.ipenz.org.nz/ipenz/Education\\_Career/accreditation/Graduate\\_Competency\\_Profiles\\_Nov\\_2009.pdf](http://www.ipenz.org.nz/ipenz/Education_Career/accreditation/Graduate_Competency_Profiles_Nov_2009.pdf).
- . 2009b. Requirements for Initial Academic Education for Professional Engineers (Version 3.1). Institute of Professional Engineers New Zealand. Available online at: [http://www.ipenz.org.nz/IPENZ/Forms/pdfs/Initial\\_Academic\\_Policy\\_Prof\\_Eng.pdf](http://www.ipenz.org.nz/IPENZ/Forms/pdfs/Initial_Academic_Policy_Prof_Eng.pdf).
- ISO. 2008. Draft International Standard ISO/DIS 31000 Risk management - Principles and guidelines on implementation (ISO/PC 992). Geneva, also available online at <http://www.iso.org>: International Organization for Standardization.
- . 2009. ISO 31000 Risk management - Principles and guidelines. Geneva, also available online at <http://www.iso.org>: International Organization for Standardization.
- Iyer, S. R., D. A. Rogers, and B. J. Simkins. 2010. Academic Research on Enterprise Risk Management. In J. Fraser and B. J. Simkins (eds), *Enterprise Risk Management: Today's Leading Research and Best Practices for Tomorrow's Executives*. The Robert W. Kolb Series in Finance, Hoboken, New Jersey: John Wiley & Sons, Inc., 419-439.
- Jackson, C. A. 2007. Strong CROs: More Important Than Ever. *Public Utilities Fortnightly* (April):56-60.
- Jackson, M. C. 2000. *Systems approaches to management*. New York: Kluwer Academic/Plenum.
- Jaeger, C. C., O. Renn, E. A. Rosa, and T. Webler. 2001. *Risk, Uncertainty, and Rational Action*. London: Earthscan Publications.
- Jaynes, E. T. 2003. *Probability Theory: The Logic of Science*, edited by G. Larry Bretthorst. Cambridge: Cambridge University Press.
- Joas, H. 1993. *Pragmatism and Social Theory*. Chicago, IL.: University of Chicago Press.
- Joerges, B. 1988. Large technical systems: Concepts and issues. In R. Mayntz and T. P. Hughes (eds), *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press, 9-36.
- Johnson, R. B., and A. J. Onwuegbuzie. 2004. Mixed Methods Research: A Research Paradigm Whose Time Has Come. *Educational Researcher* 33 (7):14-26.
- Jones, M. E. 2003. The renaissance engineer: a reality for the 21st century? *European Journal of Engineering Education* 28 (2):169-178.
- Julien, R., and L. Rieger. 2003. The Missing Link in Corporate Governance. *Risk Management* 2003 (April):32-36.

- Kahneman, D. 2003. A Perspective on Judgement and Choice: Mapping Bounded Rationality. *American Psychologist* 58 (9):697-720.
- Kaplan, S., and B. J. Garrick. 1981. On The Quantitative Definition of Risk. *Risk Analysis* 1 (1):11-27.
- Kasperson, R. E., and J. X. Kasperson. 1996. The social amplification and attenuation of risk. *Annals of the American Academy of Political and Social Science* 545:126-134.
- Kates, R. W., W. C. Clark, R. Corell, M. J. Hall, C. C. Jaeger, I. Lowe, J. J. McCarthy, H.-J. Schellnhuber, B. Bolin, N. M. Dickson, S. Faucheux, G. G. C., A. Grüber, B. Huntley, J. Jäger, N. S. Jodha, R. E. Kasperson, A. Mabogunje, P. A. Matson, H. Mooney, B. Moore III, R. O'Riordan, and U. Svedin. 2001. Sustainability Science. *Science* 292 (5517):641-642.
- Kates, R. W., and J. X. Kasperson. 1983. Comparative Risk Analysis of Technological Hazards (a review). *Proceedings of the National Academy of Sciences USA* 80:7027.
- Katko, T. S., P. S. Juuti, and P. E. Pietila. 2006. Key long-term strategic decisions in water and sanitation services management in Finland, 1860-2003. *Boreal Environment Research* 11 (2006):389-400.
- Keey, R. B. 2000. Management of engineering risk. Christchurch, NZ.: Centre for Advanced Engineering, University of Canterbury; Institution of Professional Engineers New Zealand.
- Kelly, W. J. 2007. The Risk Manager's Survival Guide. *Risk Management Magazine* 2007 (April):20-28.
- Kemp, R., S. Parto, and R. B. Gibson. 2005. Governance for sustainable development: moving from theory to practice. *International Journal of Sustainable Development* 8 (1/2):12-30.
- Kiviat, B. 2008. Reassessing Risk. *TIME* 172 (25):47-49 (22 December 2008).
- Kleffner, A. E., R. B. Lee, and B. McGannon. 2003a. The effect of corporate governance on the use of enterprise risk management: Evidence from Canada. *Risk Management and Insurance Review* 6 (1):53-73.
- . 2003b. Stronger corporate governance and its implications on risk management. *Ivey Business Journal* 67 (5):1.
- Klein, H. K., and M. D. Myers. 1999. A Set of Principles for Conducting and Evaluating Interpretive Field Studies in Information Systems. *MIS Quarterly* 23 (1):67-94.
- Klein, J. T., W. Grossenbacher-Mansuy, R. Häberli, B. Alain, R. W. Scholz, and M. Welti (eds). 2001. Transdisciplinarity: Joint Problem Solving among Science, Technology, and Society. Basel: Birkhauser Verlag.
- Kleindorfer, P. R., H. C. Kunreuther, and P. J. H. Schoemaker. 1993. Decision sciences: an integrative perspective. Cambridge, Ma.: Cambridge University Press.

- Knight, F. 1921. Risk, Uncertainty, and Profit. New York: Kelley, for the London School of Economics and Political Science.
- Koller, G. 2005. Risk Assessment and Decision Making in Business and Industry: A Practical Guide - Second Edition. Boca Raton, FL.: Chapman & Hall/CRC.
- Kovach, M. 2004. Doing justice to science: An analysis of three approaches to chemical assessment used in the negotiation of the Stockholm Convention on Persistent Organic Pollutants, Doctoral Thesis, University of California, Santa Cruz.
- KPMG International. 2007. The evolution of risk and controls: From score-keeping to strategic partnering. Published online at <http://www.us.kpmg.com>: KPMG International.
- KPMG LLP. 2009. Placing a Value on Enterprise Risk Management. Published online at <http://www.us.kpmg.com>: KPMG LLP.
- Kreiner, K. 1999. Knowledge and mind: The management of intellectual resources. In J. F. Porac and R. Garud (eds), Cognition, knowledge and organizations. *Advances in managerial cognition and organizational information processing*, 6, C. Stubbart, J. R. Meindl and J. F. Porac (eds). Stamford, CT: JAI Press Inc., 1-29.
- Krimsky, S., and D. Golding (eds). 1992. Social Theories of Risk. Westport, CT.: Praeger Publishers.
- Kuhn, T. S. 1961. The Function of Measurement in Modern Physical Science. *Isis* 52 (2):161-193.
- . 1962. The structure of scientific revolutions. Chicago, IL.: University of Chicago Press.
- La Porte, T. 1988. The United States air traffic system: Increasing reliability in the midst of rapid growth. In R. Mayntz and T. P. Hughes (eds), *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press, 215-244.
- Lafferty, W. M., and J. R. Meadowcroft. 2000. Implementing sustainable development: strategies and initiatives in high consumption societies. Oxford; New York: Oxford University Press.
- Lam, J. 2000. Enterprise-wide risk management and the role of the chief risk officer. *ERisk* (March, published online at [www.erisk.com](http://www.erisk.com)).
- . 2001. The CRO is here to stay. *Risk Management* 2001 (April):16-22.
- . 2003. Enterprise Risk Management: From Incentives to Controls. Hoboken, NJ.: John Wiley & Sons.
- . 2006. Managing Risk Across the Enterprise: Challenges and Benefits. In M. K. Ong (ed), *Risk Management: A Modern Perspective*. Burlington, MA.: Academic Press, 3-20.

- Lam, J. C., and B. M. Kawamoto. 1997. Emergence of the Chief Risk Officer. *Risk Management* 1997 (September):30-35.
- Lamsler, I., and E. Helland. 2000. How the chief risk officer is taking centre stage. *Balance Sheet* 8 (6):26-28.
- Langley, A., H. Mintzberg, P. Pitcher, E. Posada, and J. Saint-Macary. 1995. Opening up decision making: the view from the black stool. *Organization Science* 6 (3):260-279.
- Lansing, J. S. 2003. Complex Adaptive Systems. *Annual Review of Anthropology* 32 (2003):183-204.
- Laroche, H. 1995. From Decision to Action in Organisations: Decision-Making as a Social Representation. *Organization Science* 6 (1):62-75.
- Latour, B. 1987. *Science in Action: How to Follow Scientists and Engineers Through Society*. Milton Keynes: Open University Press.
- . 1988. *The Pasteurization of France*; translated by Alan Sheridan and John Law. Cambridge, MA: Harvard University Press.
- . 1996. On actor-network theory: A few clarifications. *Soziale Welt* 47:360-381.
- . 1998. From the World of Science to the World of Research? *Science* 280 (5361):208-209.
- . 1999. *Pandora's Hope: Essays on the Reality of Science Studies*. Cambridge, MA.: Harvard University Press.
- . 2005. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Oxford: Oxford University Press.
- . 2007. A Plea for Earthly Sciences: keynote lecture for the annual meeting of the British Sociological Association, East London, April 2007. Published online at [www.bruno-latour.fr/articles/article/102-BSA-GB.pdf](http://www.bruno-latour.fr/articles/article/102-BSA-GB.pdf): [Accessed 5 May 2007].
- Law, J. 1992. Notes on the Theory of the Actor Network: Ordering, Strategy and Heterogeneity. Published 2001 by the Centre for Science Studies, Lancaster University, Lancaster at [www.comp.lancs.ac.uk/sociology/papers/Law-Notes-on-ANT.pdf](http://www.comp.lancs.ac.uk/sociology/papers/Law-Notes-on-ANT.pdf): [accessed 21 April 2007].
- . 1999. After ANT: complexity, naming and topology. In J. Law and J. Hassard (eds), *Actor Network Theory and After*. Oxford: Blackwell Publishers/The Sociological Review, 1-14.
- Law, J., and V. Singleton. 2000. This is Not an Object. Published 2000 by the Centre for Science Studies, Lancaster University, Lancaster at [www.comp.lancs.ac.uk/sociology/papers/Law-Singleton-This-is-Not-an-Object.pdf](http://www.comp.lancs.ac.uk/sociology/papers/Law-Singleton-This-is-Not-an-Object.pdf): [accessed 21 April 2007].
- . 2003. *Object Lessons*. Published 2003 by the Centre for Science Studies, Lancaster

- University, Lancaster at [www.comp.lancs.ac.uk/sociology/papers/Law-Singleton-Object-Lessons.pdf](http://www.comp.lancs.ac.uk/sociology/papers/Law-Singleton-Object-Lessons.pdf). [accessed 21 April 2007].
- Law, J., and J. Urry. 2004. Enacting the social. *Economy and Society* 33 (3):390-410.
- Layton, M., and M. Fuchs. 2007. Risk Management Practices Cannot Be "Bolted On". Expert Commentary. International Risk Management Institute. Published online at <http://www.irmi.com/Expert/Articles/2007/Deloitte07.aspx>.
- Le Masurier, J., D. Blockley, and D. M. Wood. 2006. An observational model for managing risk. *Proceedings of the Institute of Civil Engineers, Civil Engineering* 159 (November):35-40.
- Lee, C. R. 2000. Chief Risk Officer Stepping Up. *Risk Management* 2000 (September):22-27.
- Lee, C. R., and P. Shimpi. 2005. The Chief Risk Officer: What Does It Look Like and How Do You Get There? *Risk Management Magazine* 2005 (September):34-38.
- Leonard, D., and W. Swap. 2004. Deep Smarts. *Harvard Business Review* September 2004:88-97.
- Levy-Leboyer, M. 1988. The French electrical power system: An inter-country comparison. In R. Mayntz and T. P. Hughes (eds), *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press, 245-262.
- Lewin, C. 2006. Enterprise risk management and civil engineering. *Proceedings of the Institute of Civil Engineers, Civil Engineering* 159 (November):4-9.
- Liebenberg, A. P., and R. E. Hoyt. 2003. The Determinants of Enterprise Risk Management: Evidence From the Appointment of Chief Risk Officers. *Risk Management and Insurance Review* 6 (1):37-52.
- Lindblom, C. E. 1959. The Science of "Muddling Through". *Public Administration Review* 19 (2):79-88.
- . 1979. Still Muddling, Not Yet Through. *Public Administration Review* 39 (6):517-526.
- Loasby, B. J. 1999. *Knowledge, Institutions and Evolution in Economics*. London: Routledge.
- Local Government Act*. (Public Act 1974 No 66). Available online at <http://www.legislation.govt.nz/act/public/1974/0066/latest/DLM415532.html>.
- Local Government Act*. (Public Act 2002 No 84). Available online at <http://www.legislation.govt.nz/act/public/1974/0066/latest/DLM415532.html>.
- Louisot, J.-P. 2003. What makes an effective Risk Manager? *Risk Management* 2003 (June):26-33.
- LTNZ. 2005. *Economic Evaluation Manual*. Wellington: Land Transport New Zealand.
- Macgill, S. M., and Y. L. Siu. 2004. The nature of risk. *Journal of Risk Research* 7 (3):315-352.
- MacGillivray, B. H., P. D. Hamilton, S. E. Hrudey, L. Reekie, and S. J. T. Pollard. 2006. Benchmarking Risk Analysis Practice in the International Water Sector. *Water*

*Practice & Technology* 1 (2):1-7.

- MacGillivray, B. H., P. D. Hamilton, J. E. Strutt, and S. J. T. Pollard. 2006. Risk Analysis Strategies in the Water Utility Sector: An Inventory of Applications for Better and More Credible Decision Making. *Critical Reviews in Environmental Science and Technology* 36 (2):85-139.
- MacGillivray, B. H., J. V. Sharp, J. E. Strutt, P. D. Hamilton, and S. J. T. Pollard. 2007a. Benchmarking Risk Management Within the International Water Utility Sector. Part I: Design of a Capability Maturity Methodology. *Journal of Risk Research* 10 (1):85-104.
- . 2007b. Benchmarking Risk Management Within the International Water Utility Sector. Part II: A Survey of Eight Water Utilities. *Journal of Risk Research* 10 (1):105-123.
- MacLean, D., R. MacIntosh, and S. Grant. 2002. Mode 2 Management Research. *British Journal of Management* 13 (2002):189-207.
- Manifest. 2004. UK Corporate Governance Milestones. Published online by Manifest, available at <http://www.manifest.co.uk/reports/index.htm>.
- March, J. G. 1978. Bounded Rationality, Ambiguity, and the Engineering of Choice. *The Bell Journal of Economics* 9 (2):587-608.
- . 1988a. *Decisions and Organizations*. Oxford: Basil Blackwell.
- . 1988b. Introduction: A chronical of speculations about organizational decision-making. In J. G. March (ed), *Decisions and Organizations*. Oxford: Basil Blackwell,
- Marsh, and RIMS. 2004. Excellence in Risk Management I: A Qualitative Survey of Successful Risk Managers. Published online at <http://www.rims.org>; Marsh and the Risk and Insurance Management Society.
- . 2005. Excellence in Risk Management II: A Qualitative Survey of Enterprise Risk Management Programs. Published online at <http://www.rims.org>; Marsh and the Risk and Insurance Management Society.
- . 2006. The Changing Face of Risk Management: Excellence in Risk Management III. Published online at <http://www.rims.org>; Marsh and the Risk and Insurance Management Society.
- . 2007. The 360° View of Risk: Excellence in Risk Management IV. Published online at <http://www.rims.org>; Marsh and the Risk and Insurance Management Society.
- . 2008. Viewing Risk Management Strategically: Excellence in Risk Management V. Published online at <http://www.rims.org>; Marsh and the Risk and Insurance Management Society.
- . 2009. Excellence in Risk Management VI: Strategic Risk Management in Practice. Published online at <http://www.rims.org>; Marsh and the Risk and Insurance Management Society.



- Marshall, R., A. Isaac, and J. Ryan. 2006. Integration of Operational Risk Management and the Sarbanes-Oxley Act Section 404. In M. K. Ong (ed), *Risk Management: A Modern Perspective*. 391-412, Burlington, MA.: Academic Press,
- Mayntz, R., and T. P. Hughes (eds). 1988. *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press.
- McDermid, D. 2006. Pragmatism. *Internet Encyclopedia of Philosophy*, eds. J. Feiser and B. Dowden: <http://www.iep.utm.edu/pragmati/>.
- McDonald, G., and M. Patterson. 2003. *Ecological Footprints of New Zealand and its Regions*. Wellington: Ministry for the Environment, Environmental Reporting Technical Paper, available online from <http://www.mfe.govt.nz/publications/ser/eco-footprint-sep03/eco-footprint-sep03.pdf>.
- McDonald, S. 2005. Studying actions in context: a qualitative shadowing method for organizational research. *Qualitative Research* 5 (4):455-473.
- McGarry, B. 2005. *Things to Think With: Understanding Interactions with Artefacts in Engineering Design*, PhD Thesis, Department of ITEE, University of Queensland, St Lucia, Brisbane, Australia.
- Megill, A. 1994. Introduction: Four Senses of Objectivity. In A. Megill (ed), *Rethinking Objectivity*. Durham and London: Duke University Press,
- Merna, T., and F. F. Al-Thani. 2005. *Corporate Risk Management: An Organisational Perspective*. Chichester, West Sussex: John Wiley and Sons.
- Meulbroek, L. K. 2002. A senior manager's guide to integrated risk management. *Journal of Applied Corporate Finance* 14 (4):56-70.
- Miccolis, J. A., K. Hively, and B. W. Merkle. 2001. *Enterprise Risk Management: Trends and Emerging Practices*. Altamonte Springs, FL.: Institute of Internal Auditors Research Foundation (prepared by Tillinghast-Towers Perrin with assistance of The Conference Board of Canada).
- Miccolis, J. A., and S. Shah. 2000. *Enterprise Risk Management: An Analytic Approach*. A Tillinghast - Towers Perrin Monograph. Published online by Towers Perrin, <http://www.towersperrin.com>.
- Mikes, A. 2005. *Enterprise risk management in action* (Discussion Paper No. 35). Published online at <http://www.lse.ac.uk/collections/CARR/documents/discussionPapers.htm>: ESRC Centre for Analysis of Risk and Regulation, London School of Economics and Political Science.
- . 2010. *Becoming the Lamp Bearer: The Emerging Roles of the Chief Risk Officer*. In J. Fraser and B. J. Simkins (eds), *Enterprise Risk Management: Today's Leading Research and Best Practices for Tomorrow's Executives*. The Robert W. Kolb Series in Finance, Hoboken, New Jersey: John Wiley & Sons, Inc., 71-86.
- Mikes, A., and D. Townsend. 2007. *Beyond Compliance: The Maturation of CROs and Other*

- Senior Risk Executives. *GARP Risk Review* 39 (November/December 2007):12-18.
- Miller, P., L. Kurunmäki, and T. O'Leary. 2006. Accounting, Hybrids and the Management of Risk. London: The London School of Economics and Political Science, Centre for Analysis of Risk and Regulation (available online at <http://www.lse.ac.uk/collections/CARR/pdf/Disspaper40.pdf>).
- Ministry of Health. 2003. Public Health Grading of Community Drinking-Water Supplies 2003: Explanatory Notes and Grading Forms. Wellington: Ministry of Health, available online at <http://www.moh.govt.nz/water>.
- . 2005. Drinking-water Standards for New Zealand (2005). Wellington: Ministry of Health, available online at <http://www.moh.govt.nz>.
- Mintzberg, H. 1970. Structured observation as a method to study managerial work. *Journal of Management Studies* 7 (1):87-104.
- Mitchell, T. 2008. Rethinking economy. *Geoforum* 39 (2008):1116-1121.
- Monstadt, J., and M. Naumann. 2005. New Geographies of Infrastructure Systems: Spatial Science Perspectives and the Socio-Technical Change of Energy and Water Supply Systems in Germany. netWORKS Papers No 10. Berlin: Leibniz-Institute for Regional Development and Structural Planning, available online at <http://www.networks-group.de/en/veroeffentlichungen/DF9163.pdf>
- . 2006. New Geographies of Infrastructure Systems: Spatial Science Perspectives and the Socio-Technical Change of Energy and Water Supply Systems in Germany. Networks Papers No 10. Berlin: Leibniz-Institute for Regional Development and Structural Planning, available online at <http://www.networks-group.de/en/veroeffentlichungen/DF9163.pdf>
- Moore, P. G., and H. Thomas. 1975. Measuring Uncertainty. *OMEGA, The International Journal of Management Science* 3 (6):657-672.
- Morgan, D. L. 2007. Paradigms Lost and Pragmatism Regained: Methodological Implications of Combining Qualitative and Quantitative Methods. *Journal of Mixed Methods Research* 1 (1):48-76.
- Morris, P. W. G., L. Crawford, D. Hodgson, M. M. Shepherd, and J. Thomas. 2006. Exploring the role of formal bodies of knowledge in defining a profession - The case of project management. *International Journal of Project Management* 24 (2006):710-721.
- Moubray, J. 1997. Reliability-Centered Maintenance, 2nd Ed. New York: Industrial Press.
- Murdoch, G. J. 1993. A History of the Water Supply of Metropolitan Auckland (preliminary draft manuscript). Auckland: Auckland Regional Authority Planning Department.
- National Academy of Engineering. 2004. The Engineer of 2020: Visions of Engineering in the New Century. Washington, D.C.: The National Academies Press.
- . 2005. Educating The Engineer of 2020: Adapting Engineering Education to the New

- Century. Washington, D.C.: The National Academies Press.
- National Commission on Fraudulent Financial Reporting. 1987. Report of the National Commission on Fraudulent Financial Reporting. Published online by the Committee of Sponsoring Organisations of the Treadway Commission (COSO), available at <http://www.coso.org/Publications/NCFFR.pdf>.
- NIST. 2006. Criteria for Performance Excellence. Gaithersburg, MD.: Baldrige National Quality Program, National Institute of Standards and Technology, Technology Administration, Department of Commerce (<http://www.quality.nist.gov>).
- Nocco, B. W., and R. M. Stulz. 2006. Enterprise Risk Management: Theory and Practice. *Journal of Applied Corporate Finance* 18 (4):8-20.
- Noël, A. 1989. Strategic Cores and Magnificent Obsessions: Discovering Strategy Formation Through Daily Activities of CEOs. *Strategic Management Journal* 10:33-49.
- Nonaka, I. 1994. A Dynamic Theory of Organizational Knowledge Creation. *Organization Science* 5 (1):14-37.
- Nonaka, I., and N. Konno. 1998. The concept of "Ba": Building a foundation for knowledge creation. *California Management Review* 40 (3):40-54.
- Nonaka, I., and H. Takeuchi. 1995. *The Knowledge Creating Company*. London: Oxford University Press.
- Nowotny, H., P. Scott, and M. Gibbons. 2001. *Re-Thinking Science: Knowledge and the Public in an Age of Uncertainty*. Cambridge; Malden, MA: Polity; Blackwell.
- . 2003. Introduction: 'Mode 2' Revisited: The New Production of Knowledge. *Minerva* 41 (2003):179-194.
- O'Rourke, M. 2007. Scaling the heights of ERM. *Risk Management Magazine* 2007 (October, Spotlight Issue):52-56.
- objectivity. 2000a. In *The Blackwell Dictionary of Sociology*. Oxford: Blackwell Publishers. Retrieved January 14, 2009, from <http://www.credoreference.com.ezproxy.auckland.ac.nz/entry/724180/>.
- . 2000b. In *The Dictionary of Human Geography*. Oxford: Blackwell Publishers. Retrieved January 14, 2009, from <http://www.credoreference.com.ezproxy.auckland.ac.nz/entry/734499/>.
- OCEG. 2007. OCEG Matrix Adapted From RIMS ERM Risk Maturity Model. Available online at <http://www.oceg.org/Details/RIMSERM: Open Compliance and Ethics Group>.
- Osbourne, T., and N. Rose. 1999. Do the social sciences create phenomena? The example of public opinion research. *British Journal of Sociology* 50:367-396.
- Pagach, D., and R. Warr. 2008a. The Characteristics of Firms that Hire Chief Risk Officers. North Carolina State University Working Paper (May 2008). Available at

[http://www.mgt.ncsu.edu/erm/documents/Risk\\_officer\\_hazard\\_JRI2.pdf](http://www.mgt.ncsu.edu/erm/documents/Risk_officer_hazard_JRI2.pdf).

- . 2008b. The effects of enterprise risk management on firm performance. North Carolina State University Working Paper. Available at [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1155218](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1155218).
- Pahl, G., W. Beitz, J. Feldhusen, and K.-H. Grote. 2007. *Engineering Design: A Systematic Approach*, translated and edited by K. Wallace and L. Blessing, 3rd ed. London: Springer-Verlag.
- Pandza, K., and R. Thorpe. 2010. Management as Design, but What Kind of Design? An Appraisal of the Design Science Analogy for Management. *British Journal of Management* 21 (2010):171-186.
- Paskiewicz, L. S. 2002. The Shadowing Experience: Valuing the Link Between Faculty Practice and Student Learning. *Journal of Professional Nursing* 18 (4):238-42.
- Peet, J. 2000. Being Fully Human and Creating a Better Future: Sustainable Development from an Integrated Systems Perspective. Baltic21 Workshop on Sustainable Development, 8-9 June, Sigtuna, Sweden.
- Pepper, S. C. 1942. *World hypotheses: a study of evidence*. Los Angeles and London: University of California Press.
- Perlow, L. A. 1998. Boundary Control: The social Ordering of Work and Family Time in a High-Tech Corporation. *Administrative Science Quarterly* 43 (2):328-357.
- Petit, C. 2006. The clairvoyant CRO. Published online at <http://www-03.ibm.com/industries/financialservices/doc/content/resource/thought/1594875103.html>: IBM Business Consulting Services.
- Pickering, A. 1993. The Mangle of Practice: Agency and Emergence in the Sociology of Science. *The American Journal of Sociology* 99 (3):559-589.
- . 1994. Objectivity and the Mangle of Practice. In A. Megill (ed), *Rethinking Objectivity*. Durham and London: Duke University Press,
- Polanyi, M. 1962. *Personal Knowledge*. Chicago, IL.: University of Chicago Press.
- Polite, V. C., R. McClure, and D. L. Rollie. 1997. The Emerging Reflective Urban Principal - The Role Of Shadowing Encounters. *Urban Education* 31 (5):466-489.
- Pollard, S. J. T., J. E. Strutt, B. H. MacGillivray, P. D. Hamilton, and S. E. Hrudey. 2004. Risk analysis and management in the water utility sector: A review of drivers, tools, and techniques. *Trans. IChemE, Part B: Process Safety and Environmental Protection* 82 (B6):453-462.
- Porter, L. W., and L. E. McKibbin. 1988. *Management Education and Development*. New York: McGraw-Hill.
- Porter, T. M. 1995. Trust in numbers: the pursuit of objectivity in science and public life.

- Princeton: Princeton University Press.
- Power, M. 2004. *The Risk Management of Everything: Rethinking the politics of uncertainty*. London: Demos (available online at [www.demos.co.uk](http://www.demos.co.uk)).
- . 2005a. Enterprise Risk Management and Organization of Uncertainty in Financial Institutions. In K. Knorr-Cetina and A. Preda (eds), *The Sociology of Financial Markets*. Oxford: Oxford University Press, 250-268.
- . 2005b. Organizational responses to risk: the rise of the chief risk officer. In B. Hutter and M. Power (eds), *Organizational Encounters with Risk*. Cambridge: Cambridge University Press, 132-148.
- . 2007. *Organized Uncertainty*. Oxford: Oxford University Press.
- PriceWaterhouseCoopers. 2002. *Taming Uncertainty: Risk management the entire enterprise*. Global Financial Services e-briefing programme. Published online at [www.pwc.com/en\\_GX/gx/financial-services/.../111802taming.pdf](http://www.pwc.com/en_GX/gx/financial-services/.../111802taming.pdf).
- . 2007. Independent review of WSL's risk management framework (June 2007). Report to Watercare Services Ltd.
- . 2008. Does ERM Matter? Enterprise risk management in the insurance industry, a global study. Published online at [www.pwc.com/en\\_GX/gx/insurance/pdf/erm\\_survey.pdf](http://www.pwc.com/en_GX/gx/insurance/pdf/erm_survey.pdf) (June 2008).
- . 2009. *Seizing Opportunity: Linking risk and performance*. Thought Leadership Institute. Published online at [www.pwc.com/us/en/risk.../linking-risk-peformance-download.jhtml](http://www.pwc.com/us/en/risk.../linking-risk-peformance-download.jhtml) (July 2009).
- Quinn, L. R. 2004. The CRO: A Species at Risk. *Risk & Insurance* 15 (10):20-22.
- Ramsey, F. P. 1926. Truth and Probability. In Ramsey (1931) *The Foundations of Mathematics and other Logical Essays*, Ch. VII, p 156-198, R. B. Braithwaite (Ed.). London: Kegan, Paul, Trench, Trubner & Co., New York: Harcourt, Brace and Company. (1999 electronic version, History of Economic Thought Website, <http://homepage.newschool.edu/het//texts/ramsey/ramsess.pdf>, downloaded 7 September 2010).
- Rechar, R. P. 1999. Historical Relationship Between Performance Assessment for Radioactive Waste Disposal and Other Types of Risk Assessment. *Risk Analysis* 19 (5):763-807.
- Renn, O. 1992. Concepts of Risk: A Classification. In S. Krimsky and D. Golding (eds), *Social Theories of Risk*. Westport, CT.: Praeger Publishers, 53-82.
- . 1998. Three decades of risk research: accomplishments and new challenges. *Journal of Risk Research* 1 (1):49-71.
- Rescher, N. 1983. *Risk: A Philosophical Introduction to the Theory of Risk Evaluation and Management*. Washington, D.C.: University Press of America.

- . 1995. Satisfying Reason: Studies in the Theory of Knowledge. *Episteme*, 21, M. Bunge (ed). Dordrecht, The Netherlands: Kluwer Academic Publishers.
- . 2005. Realism and Pragmatic Epistemology. Pittsburgh, PA.: University of Pittsburgh Press.
- RIMS. 2006. Risk Maturity Model for Enterprise Risk Management. Online at <http://www.rims.org/resources/ERM/Pages/RiskMaturityModel.aspx>.: Risk and Insurance Management Society, Inc.
- risk. 2000. In Collins English Dictionary. Retrieved 20 February 2008, from <http://www.credoreference.com.ezproxy.auckland.ac.nz/entry/2687667risk>.
- . 2008. In Merriam-Webster Online. Retrieved 20 February 2008, from <http://www.merriam-webster.com/dictionary/risk>.
- Rittel, H., and M. Webber. 1973. Dilemmas in a General Theory of Planning. *Policy Sciences* 4 (1973):155-169.
- RMRDPC. 2002. Risk Management Maturity Level Development. Risk Management Research and Development Program Collaboration, available online at <http://www.pmi-switzerland.ch/fall05/riskmm.pdf>.
- Robbins, P. T. 2007. The reflexive engineer: Perceptions of integrated development. *Journal of International Development* 19 (2007):99-110.
- Robèrt, K.-H., B. Schmidt-Bleek, J. Aloisi de Larderel, G. Basile, J. L. Jansen, R. Kuehr, P. Price Thomas, M. Suzuki, P. Hawken, and M. Wackernagel. 2002. Strategic sustainable development -- selection, design and synergies of applied tools. *Journal of Cleaner Production* 10 (3):197-214.
- Roberts, D. 2006. The Chief Risk Officer: Finding New Ways Forward. *The RMA Journal* 88 (11):66-69.
- Roberts, N. 2001. Coping with wicked problems: the case of Afghanistan. In L. Jones, J. Guthrie and P. Steane (eds), Learning from International Public Management Reform. Research in Public Policy Analysis and Reform, Vol. 11B, Greenwich, Conn.: JAI Press, 353-376.
- Robinson, J. 2004. Squaring the circle? Some thoughts on the idea of sustainable development. *Ecological Economics* 48 (2004):369-384.
- Rochlin, G. I. 1994. Broken Plowshare: System Failure and the Nuclear Power Industry. In J. Summerton (ed), Changing Large Technical Systems. Boulder, Co.: Westview Press, 231-264.
- Rosa, E. A. 1998. Metatheoretical foundations for post-normal risk. *Journal of Risk Research* 1 (1):15-44.
- Roth, J., L. Berg, and A. Styhre. 2004. Knowledge Facilitation in Action. In N. Adler, A. B. Shani and A. Styhre (eds), Collaborative Research in Organizations: Foundations for

- Learning, Change, and Theoretical Development. Thousand Oaks, Ca.: Sage Publications,
- Rotmans, J., R. Kemp, and M. van Asselt. 2001. More evolution than revolution: transition management in public policy. *Foresight* 3 (1):15-31.
- Runde, J. 1998. Clarifying Frank Knight's discussion of the meaning of risk and uncertainty. *Cambridge Journal of Economics* 22 (1998):539-546.
- Ryle, G. 1949. *The Concept of Mind*. London: Hutchinson & Co.
- Salsbury, S. 1988. The emergence of an early large-scale technical system: The American railroad network. In R. Mayntz and T. P. Hughes (eds), *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press, 37-68.
- Salvador, S. 2007. Risk Management Governance: Creating the Foundation for ERM. *The RMA Journal* 89 (5):30-33.
- Schaefer, D., J. H. Panchal, S.-K. Choi, and F. Mistree. 2008. Strategic Design of Engineering Education for the Flat World. *International Journal of Engineering Education* 24 (2):274-282.
- Schanfield, A., and D. Helming. 2008. 12 Top ERM Implementation Challenges. *Internal Auditor* 65 (6):41-44.
- Scheffer, M., S. Carpenter, J. A. Foley, C. Folke, and B. Walker. 2001. Catastrophic shifts in ecosystems. *Nature* 413 (11 October):591-596.
- Schermerhorn, J. R. 1996. *Management*. 5th ed. New York: John Wiley & Sons.
- Schön, D. 1983. *The reflective practitioner: how professionals think in action*. New York: Basic Books.
- Schrader-Frechette, K. 1991. *Risk and Rationality: Philosophical Foundations for Populist Reforms*. Berkeley and Los Angeles: University of California Press.
- Schwitzgebel, E. 2010. Belief. *The Stanford Encyclopedia of Philosophy* (Summer 2010 Edition), ed. E. N. Zalta: <http://plato.stanford.edu/archives/sum2010/entries/belief/>.
- Scott, S. V., and N. Perry. 2006. The enactment of risk categories: Organizing and re-organizing risk management practices in the energy industry (Working Paper Series, No. 148). Published online at <http://www.lse.ac.uk/>: Department of Management Information Systems Group, London School of Economics and Political Science.
- SEC. 2009. SEC Approves Enhanced Disclosure About Risk, Compensation and Corporate Governance. US Securities and Exchange Commission, Press Release 2009-268, 16 December. Retrieved 20 January 2010 from <http://www.sec.gov/news/press/2009/2009-268.htm>.
- Seely, B. E. 1999. The other re-engineering of engineering education, 1900-1965. *Journal of*

*Engineering Education* 88 (3):285-294.

- . 2005. Patterns in the History of Engineering Education Reform: A Brief Essay. In National Academy of Engineering (ed), *Educating the Engineer of 2020: Adapting Engineering Education to the New Century*. Washington, D.C.: The National Academies Press, 114-130.
- Shafir, E., and R. A. LeBoeuf. 2002. Rationality. *Annual Review of Psychology* 53 (2002):491-517.
- Shani, A. B., S. A. Mohrman, W. A. Pasmore, B. Stymne, and N. Adler (eds). 2008. *Handbook of Collaborative Management Research*. Thousand Oaks, CA.: Sage.
- Sheffield, J. 2005. Systemic Knowledge and the V-Model. *International Journal of Business Information Systems* 1 (1/2):83-101.
- Sheffield, J., and R. Gou. 2007. Ethical Inquiry in Knowledge Management. *International Journal of Applied Systemic Studies* 1 (1):68-81.
- Shenkir, W. G., and P. L. Walker. 2006. Implementing Enterprise Risk Management. Draft Statement on Management Accounting. Published online at [www.imanet.org/pdf/DraftSMAonERM.pdf](http://www.imanet.org/pdf/DraftSMAonERM.pdf): Institute of Management Accountants. .
- . 2007. Enterprise Risk Management: Tools and Techniques for Effective Implementation. Statement on Management Accounting. Published online at <http://www.imanet.org>: Institute of Management Accountants.
- Short, J. F. 1992. Defining, Explaining, and Managing Risks. In J. F. Short and L. Clarke (eds), *Organizations, Uncertainties, and Risks*. Boulder, CO.: Westview Press, 3-23.
- Short, J. F., and L. Clarke (eds). 1992. *Organizations, Uncertainties, and Risk*. Boulder, CO.: Westview Press.
- Simon, H. A. 1955. A behavioural model of rational choice. *Quarterly Journal of Economics* 69 (1):99-118.
- . 1957. *Administrative Behaviour: A Study of Decision-Making Processes in Administrative Organizations* (2nd Ed.). New York: Macmillan.
- . 1960. *The new science of management decision*. New York: Harper.
- Smith, A. M. 1993. *Reliability Centered Maintenance*. New York: McGraw-Hill.
- Sneddon, C., R. B. Howarth, and R. B. Norgaard. 2006. Sustainable development in a post-Brundtland world. *Ecological Economics* 57 (2006):253-268.
- Snowden, D. 2002. Complex acts of knowing: Paradox and descriptive self-awareness. *Journal of Knowledge Management* 6 (2):100-111.
- Snyder, N., and W. F. Glueck. 1980. How Managers Plan - The Analysis of Managers' Activities. *Long Range Planning* 13:70-6.



- Spencer Brown, G. 1969. *The Laws of Form*. London: Allen & Unwin.
- Spender, J.-C., and A. G. Scherer. 2007. The Philosophical Foundations of Knowledge Management. *Organization* 14- (1):5-28.
- Spinks, N., N. L. J. Silburn, and D. W. Birchall. 2007. Making it all work: the engineering graduate of the future, a UK perspective. *European Journal of Engineering Education* 32 (3):325-335.
- Standards Australia International Limited, and Standards New Zealand. 2004a. Australia/New Zealand Standard 4360: Risk Management. Sydney / Wellington: Standards Australia International / Standards New Zealand.
- . 2004b. HB 436:2004 Risk Management Guidelines Companion to AS/NZS 4360:2004. Sydney / Wellington: Standards Australia International / Standards New Zealand.
- Stanley, J., and T. Williamson. 2001. Knowing How. *The Journal of Philosophy* 98 (8):411-444.
- Starkey, K., A. Hatchuel, and S. Tempest. 2009. Management Research and the New Logics of Discovery and Engagement. *Journal of Management Studies* 46 (3):547-558.
- Starkey, K., and P. Madan. 2001. Bridging the Relevance Gap: Aligning Stakeholders in the Future of Management Research. *British Journal of Management* 12 (Special Issue):S3-S26.
- Statistics New Zealand. 2007. Subnational Population Estimates. Table 2 Estimated Resident Population, Territorial Authorities, as at 30 June 2001-2007. Published online at <http://www.stats.govt.nz/tables/subnat-pop-estimates-tables.htm> [SubPopEstJune07alltables.xls].
- Steffen, W., A. Sanderson, J. Jäger, P. D. Tyson, B. Moore III, P. A. Matson, K. Richardson, F. Oldfield, H.-J. Schellnhuber, B. J. Turner II, and R. J. Wasson. 2004. *Global Change and the Earth System: A Planet Under Pressure*. Heidelberg, Germany: International Geosphere-Biosphere Programme Book Series, Springer Verlag.
- Steup, M. 2010. Epistemology. *The Stanford Encyclopedia of Philosophy* (Spring 2010 Edition), ed. E. N. Zalta: <http://plato.stanford.edu/archives/spr2010/entries/epistemology/>.
- Stroh, P. J. 2005. Enterprise risk management at United Healthcare. *Strategic Finance* 87 (1):24-35.
- Stulz, R. M. 2009. 6 Ways Companies Mismanage Risk. *Harvard Business Review* 87 (3):86-94.
- Sydenham, P. H. 2004. *Systems approach to engineering design*. Norwood, MA: Artech House.
- Tainter, J. A. 1988. *The Collapse of Complex Societies*. Cambridge: Cambridge University Press.
- Taleb, N. N., D. G. Goldstein, and M. W. Spitznagel. 2009. The six mistakes executives make in risk management. *Harvard Business Review* 87 (10):78-81.

- Tarantino, A. 2006. *Manager's Guide to Compliance: Best Practices and Case Studies*. Hoboken, NJ.: John Wiley & Sons.
- The Economist. 2009. Silo but deadly. Issue No. 8660, December 5th-11th, 73-74.
- The Treasury. 2008. Demonstrating Value-for-Money. <http://www.treasury.govt.nz/publications/guidance/performance>. [Accessed 17 March 2009].
- Thomas, F. 1988. The politics of growth: The German telephone system. In R. Mayntz and T. P. Hughes (eds), *The Development of Large Technical Systems*. Frankfurt am Main: Campus Verlag; and Boulder, Co.: Westview Press, 179-214.
- Thompson, P. 1986. The Philosophical Foundations of Risk. *Southern Journal of Philosophy* 24 (2):273-286.
- Towers Perrin. 2006. A Changing Risk Landscape: A Study of Corporate ERM in the U.S. Published online at <http://www.towersperrin.com>: Towers Perrin.
- Towers Perrin Tillinghast. 2006. Risk Management. Risk Opportunity. The 2006 Tillinghast ERM Survey. Published online at <http://www.towersperrin.com>: Towers Perrin Tillinghast.
- Tranfield, D., and K. Starkey. 1998. The Nature, Social Organization and Promotion of Management Research: Towards Policy. *British Journal of Management* 9 (1998):341-353.
- Transfund New Zealand. 1997. *Project Evaluation Manual*. Wellington: Transfund New Zealand.
- Tsoukas, H. 2003. Do We Really Understand Tacit Knowledge. In M. Easterby-Smith and M. A. Lyles (eds), *The Blackwell Handbook of Organizational Learning and Knowledge Management*. Oxford: Blackwell Publishing, 410-427.
- Tsoukas, H., and E. Vladimirou. 2001. What is organizational knowledge? *Journal of Management Studies* 38 (7):973-993.
- Tuohy, C. 2006. CRO: The Kiss of Death. *Risk & Insurance* 17 (9):50.
- UKWIR. 2002. *Capital Maintenance Planning: A Common Framework, Vols I - IV*. Report No. 02/RG/05/3. UK Water Industry Research Limited (<http://www.ukwir.org>).
- UNEP et al. 1997. *Engineering Education and Training for Sustainable Development*. Report of the conference jointly sponsored by the United Nations Environment Program (UNEP), the World Federation of Engineering Organisations (WFEO), the World Business Council for Sustainable Development (WBCSD), and the Ecole Nationale des Ponts et Chaussées (ENPC), 24-26 September in Paris, France. Available online at <http://www.iies.es/FMOI-WFEO/desarrollosostenible/main/assets/ParisReport.doc>.
- Utter, J. S. 2006. Enterprise Risk Management in 2005 - Moving Beyond Market and Credit Risk. In K. Leggio, D. L. Bodde and M. L. Taylor (eds), *Managing Enterprise Risk: What*

- The Electric Industry Experience Implies For Contemporary Business. Oxford, UK.: Elsevier, 107-115.
- van Aken, J. E. 2005. Management Research as Design Science: Articulating the Research Products of Mode 2 Knowledge Production in Management. *British Journal of Management* 16 (2005):19-36.
- van Bueren, E. M., E.-H. Klijn, and J. F. M. Koppenjan. 2003. Dealing with Wicked Problems in Networks: Analyzing an Environmental Debate from a Network Perspective. *Journal of Public Administration Research and Theory* 13 (2):193-212.
- van der Brugge, R., J. Rotmans, and D. Loorbach. 2005. The transition in Dutch water management. *Regional Environmental Change* 5 (4):164-176.
- van der Vleuten, E. 2004. Infrastructures and Societal Change. A View from the Large Technical Systems Field. *Technology Analysis and Strategic Management* 16 (3):395-414.
- Vinten, G. 2001. Corporate Governance and the Sons of Cadbury. *Corporate Governance* 1 (4):4-8.
- von Krogh, G., K. Ichijo, and I. Nonaka. 2000. Enabling Knowledge Creation. New York: Oxford University Press.
- Voß, J.-P., and R. Kemp. 2006. Sustainability and reflexive governance: introduction. In J.-P. Voß, D. Bauknecht and R. Kemp (eds), Reflexive governance for sustainable development. Cheltenham: Edward Elgar, 3-32.
- Vukic, A., and B. Keddy. 2002. Northern Nursing Practice in a Primary Health Care Setting. *Journal of Advanced Nursing* 40 (5):542-548.
- Waldrop, M. M. 1992. Complexity. New York: Simon & Schuster.
- Walker, C. R., R. H. Guest, and A. N. Turner. 1956. The Foreman on the Assembly Line. Boston, Ma.: Harvard University Press.
- Walker, P. L., and W. G. Shenkir. 2008. Implementing Enterprise Risk Management. *Journal of Accountancy* 205 (3):31.
- Wallace, R. J. 2009. Practical Reason. The Stanford Encyclopedia of Philosophy (Summer 2009 Edition), ed. E. N. Zalta: <http://plato.stanford.edu/entries/practical-reason/>.
- Wan, X., H. Nakatani, K. Ueno, T. Asamizuya, K. Cheng, and K. Tanaka. 2011. The Neural Basis of Intuitive Best Next-move Generation in Board Game Experts. *Science* 331 (6015):341-346.
- Ward, J. 2006. ERM: The New Standard And Practice In Good Corporate Governance. *RatingsDirect* 2006 (June 14):1-8, published by Standard & Poors, available online at <http://www2.standardandpoors.com>.
- Ward, S. 2001. Exploring the Role of the Corporate Risk Manager. *Risk Management: An*

*International Journal* 3 (1):7-25.

———. 2003. Approaches to Integrated Risk Management: A Multi-Dimensional Framework. *Risk Management* 5 (4):7-23.

———. 2005. Risk management, organisation and context. London: Witherbys & Co. Ltd, for the Institute of Risk Management, London.

Waring, A., and A. I. Glendon. 1998. *Managing Risk: Critical issues for survival and success into the 21st Century*. London: International Thomson Business Press.

Watercare Services Ltd. 2000. Annual Report. Auckland: Watercare Services Ltd.

———. 2002a. Annual Report. Auckland: Watercare Services Ltd.

———. 2002b. Wastewater Treatment: The Mangere Wastewater Treatment Plant. Auckland: Watercare Services Ltd.

———. 2003a. Annual Report. Auckland: Watercare Services Ltd.

———. 2003b. Information Sheet: The History of Wastewater Treatment in Auckland. Available online at <http://www.watercare.co.nz/assets/Wastewater/WWHistory.pdf>.

———. 2003c. Risk Management Framework. Auckland, New Zealand.

———. 2004a. Annual Report. Auckland: Watercare Services Ltd.

———. 2004b. Constitution of Watercare Services Ltd.

———. 2005a. Annual Report. Auckland: Watercare Services Ltd.

———. 2005b. Asset Management Plan 2007-2026. Auckland: Watercare Services Ltd.

———. 2005c. Statement of Corporate Intent (for the three year period commencing 1 July 2005). Auckland: Watercare Services Ltd.

———. 2005d. A Three Waters Vision. Auckland: Auckland Water Industry, facilitated by Watercare Services Ltd.

———. 2006. Asset Management Plan 2008-2027. Auckland: Watercare Services Ltd.

———. 2007a. Annual Report. Auckland: Watercare Services Ltd.

———. 2007b. Asset Management Plan 2009-2028. Auckland: Watercare Services Ltd.

———. 2008a. Asset Management Plan 2010-2029. Auckland: Watercare Services Ltd.

———. 2008b. Three Waters Strategic Plan (discussion version). Auckland: Auckland Water Industry, facilitated by Watercare Services Ltd.

Watson, T. J. 1997. Theorizing Managerial Work: a Pragmatic Pluralist Approach to Interdisciplinary Research. *British Journal of Management* 8 (1997):3-8.

- . 2011. Ethnography, Reality, and Truth: The Vital Need for Studies of 'How Things Work' in Organizations and Management. *Journal of Management Studies* 48 (1):202-217.
- Weait, M. 1993. Icing on the cake? The contribution of the compliance function to effective financial regulation. *Journal of Asset Protection and Financial Crime* 1 (1):83-90.
- Weaver, P., L. Jansen, G. van Grootveld, E. van Spiegel, and P. Vergragt. 2000. Sustainable Technology Development. Sheffield: Greenleaf Publishing.
- Weber, E. U., and E. J. Johnson. 2009. Mindful Judgement and Decision Making. *Annual Review of Psychology* 60 (2009):53-85.
- Wesner, J. W., and C. L. Dym. 2008. What we have learned at Mudd Design Workshop VI, 'Design Engineering Education in a Flat World'. *International Journal of Engineering Education* 24 (2):443-448.
- WFE0. 1991. The 'Arusha Declaration'. Statement by the World Federation of Engineering Organisations to the United Nations Conference on Environment and Development 1992. Adopted at the WFE0 General Assembly, Arusha, Tanzania, September 1991. Available at [http://www.wfeo.org/documents/download/arusha\\_declaration.pdf](http://www.wfeo.org/documents/download/arusha_declaration.pdf).
- . 1997. The Engineer's Response to Sustainable Development. World Federation of Engineering Organizations.
- . 2004. The Shanghai Declaration on Engineering and the Sustainable Future. Declaration of the World Engineers' Convention 2004, sponsored by the World Federation of Engineering Organizations (WFE0) and co-sponsored by the United Nations Educational, Scientific and Cultural Organization (UNESCO), and organized by the China Association for Science and Technology (CAST), Chinese Academy of Engineering (CAE), and Shanghai Municipal Government, 2-6 November, Shanghai, China. Available at [http://www.wfeo.org/documents/download/WEC2004\\_Shanghai\\_%20Final%20Report.pdf](http://www.wfeo.org/documents/download/WEC2004_Shanghai_%20Final%20Report.pdf).
- Wheeler, J. A. 2009. The Rise of the Chief Risk Officer. *Internal Auditor* 65 (3):55-57.
- Whitcomb, J. 2010. SEC Requires ERM as of Feb 28 2010. Promotional e-mail from LogicManager Inc. to RIMS mailing list. Received 20 January 2010.
- Witte, E. 1972. Field research on complex decision making processes - the phase theorem. *International Studies of Management and Organization* 2 (2):156-182.
- Wolcott, H. F. 2003. The Man in the Principal's Office. Walnut Creek, Ca.: Altamira Press.
- Wood, D. 2002. From cop to CRO. *ERisk* (March, published online at [www.erisk.com](http://www.erisk.com)).
- Wood, M. 2008. Using numbers in risk management - pitfalls and practicalities. 6th Annual Enterprise Risk Management Conference, 28 - 29 April 2008, Auckland, New Zealand. Conferenz Ltd.

- Woodrow, M. 2005. What makes an exceptional risk manager? *Risk Management* 52 (5):41.
- World Commission on Environment and Development. 1987. *Our Common Future*. Oxford: Oxford University Press.
- Wright, J. F. 2002. *Monte Carlo Risk Analysis and Due Diligence of New Business Ventures*. New York: Amacom.
- WSAA. 2003. *Asset Management Benchmarking Framework*: Water Services Association of Australia.
- . 2008. *Aquamark Users Manual (January 2008)*: Water Services Association of Australia.
- Yankelovich, D. 1999. *The magic of dialogue: transforming conflict into cooperation*. New York: Simon & Schuster.
- Yin, L., M. A. J. Smith, and K. S. Trivedi. 2001. *Uncertainty Analysis in Reliability Modeling*. Annual Reliability and Maintainability Symposium, 22 - 25 January 2001, Philadelphia, PA.
- Zierhofer, W., and P. Burger. 2007. Disentangling Transdisciplinarity: An Analysis of Knowledge Integration in Problem-Oriented Research. *Science Studies* 20 (1):51-74.