Knowledge Management Research & Practice (2003) 1, 113–119

www.palgrave-journals.com/kmrp

POSITION PAPER

Innovation as an objective of knowledge management. Part I: The landscape of management

Dave Snowden[1]

[1]*Canolfan Cynefin Centre, IBM Global Services, UK*

Correspondence: Dave Snowden, Canolfan Cynefin Centre, IBM Global Services, University of Cardiff, UK.
E-mail: snowded@uk.ibm.com

Abstract

This is the first of a two-part position paper considering how knowledge management supports innovation. The focus of the first part is on diagnosis, and the development of management science and management practice in the 20th century. It is argued that systems need to be understood not as systems, but at the level of agent interaction, and that an emphasis on design should give way to an emphasis on emergence. The second part will develop these ideas further, with a focus on intervention.

Knowledge Management Research & Practice (2003) **1,** 113–119. doi:10.1057/palgrave.kmrp.8500014

Keywords: complexity; emergence; innovation; knowledge management theory; systems

Received: 20 October 2003

Introduction

Creating the conditions for innovation, along with enabling better decision making are to my mind the primary goals or functions of Knowledge Management (KM). In *Complex Acts of Knowing* (Snowden, 2002), I argued that the post-1995 focus on tacit to explicit knowledge conversion is flawed, in that it sees tacit and explicit knowledge as distinct categories rather than perspectives or aspects of a knowledge act. Instantiated in the SECI model (Nonaka & Takeuchi, 1995) the distinction failed in practice to recognise that *we always know more than we can say, and we will always say more than we can write down.* Recognition of this principle led me to argue that we need to separate KM into the distinct management of context, narrative and content, and that the current emphasis on content management was an inevitable consequence of misinterpretation of Polanyi's (1962) distinction between tacit and explicit knowledge in the use (and abuse) of the SECI model. More recent attempts in the first edition of this journal to mitigate the impact of tacit-explicit knowledge conversation using a Western concept of dialectic (Nonaka & Toyama, 2003) are to my mind still incomplete in that they (i) confuse dichotomy and dilemma with dialectic, seeing dialectic as a form of synthesis of existing understandings rather than as a means of evolving a different level understanding and (ii) they, to use a complexity concept, are subject to premature convergence to a categorisation model of knowledge without first going through the necessary sense-making phase. The meaning of these two somewhat cryptic points should be evident by the end of this paper.

Moving beyond criticism of the dominant practice of KM, *Complex Acts of Knowing* utilised a growing understanding of the application of complexity science to create a sense-making framework based on different underlying relationships between cause and effect. The framework was then used to determine different types of KM practice. The concept of contextual complexity (now more fully developed in this article as contingent complexity) as represented in the Cynefin sense-making framework, has been applied to decision making (Kurtz & Snowden, 2003) and I now continue my consideration of the post-SECI generation of KM by applying the framework to innovation with a result that somewhat resembles the bride's nursery rhyme: *something old, something new, something borrowed, something blue*.

This paper consists of two parts, the first is diagnosis and the second is intervention. In the first part, I want to look at the development of management science and management practice over the past century, arguing that both have been dominated by approaches based on an underlying belief that systems, by which I mean markets, organisations, etc. are ordered with discoverable and repeating relations between cause and effect. Apparently opposing approaches such as Business Process Engineering (Hammer & Champy, 1993) and the Learning Organisation (Senge, 1990) both assume that systems can be understood as a whole and designed as such. I will contrast those design-based approaches with new concepts from complexity science. This applies to systems that cannot be understood at the level of the whole system, other than in retrospect, but have to be understood at the level of agent interaction from which order emerges. My overall argument is that the systems level design approaches (abbreviated to design hereafter) dominate KM practice, which encompasses both engineering (often expressed as technology projects) and learning (often expressed in various guises as issues relating to culture). I further argue that design is now given in KM and there is no excuse not to do it well, but that design alone no longer offers either competitive advantage or novelty, in service provision. For advantage and novelty we have to look at systems that are emergent, where what design that does take place is design at the level of agent interaction and cannot to based on strategic goals, gap analysis and input–output-type models.

In Part II, I will update and summarise the language of the Cynefin sense-making framework that presumes an ontological separation of domains based on different types of causality. A fuller elaboration of this is provided in *The New Dynamics of Strategy* (Kurtz & Snowden, 2003), but familiarity with that material is not assumed. I will then look at four of what are currently 18 different dynamic movements or 'epistemological trajectories' across those domains that provide for innovation within organisations and then conclude by addressing innovative capacity as a problem of KM through the filters of context, narrative and content.

Concept and practice

Method development in management consultancy practice (and regrettably in far too much academic work in the field) normally involves gathering data from organisations, the derivation from that data of a hypothesis that can be tested against other data to create a predictive or prescriptive recipe for use in organisations. The hypothesis arises from a series of cases and suggests that if these five, 10, 50 or whatever organisations that are deemed to be successful adopt this particular management practice, then the adoption of that management practice will result in success. In effect, this is a confusion of properties and qualities that is satirised in first-year Philosophy courses, by pointing out that just because I have met a French person wearing glasses it does not follow that all French people wear glasses, or that wearing glasses is an essential feature of being French. The psychology of executives who seek such simple prescriptions despite the fact that few successful companies have ever innovated by imitating past best practice is beyond the scope of this paper but deserves some study; possibly an opportunity for cross-disciplinary study with a specialist in the behaviour of lemmings.

This approach has been challenged recently in a well-constructed article in the Harvard Business review (Christensen & Raynor, 2003). It argues that executives need to pay more attention to management theory rather than paying attention to simple recipes derived from a superficial understanding of past practice in other organisations 'in the naïve belief that if a particular course of action helped other companies to succeed, it ought to help theirs too.' The article states that good theory is held to be part of a three-stage process with commences with the observation and description of the phenomenon we wish to explore, which are then sorted into categories from which a theory is formulated. The process is then cyclical as the theory is used to predict future observations, with the theory being refined or refuted based on the results of those observations.

The excellence of the article lies in its identification of the dangers of using a limited amount of data to construct general theories, the way in which correlations are used to 'masquerade as causation' and the frequent use of partial selections of data from the same organisations to prove contradictory theories. Hopefully, this HBR article will encourage executives to avoid simple recipes or better still discourage them from the all to frequent anti-theory approach in which 'academic' is considered a term of abuse. However, it does not challenge the need for recipes or the presumption that causality can be established through categorisation. In this article, I wish to argue for (i) a dialectical relationship between categorisation and sense making in the context of innovation within organisations, (ii) a legitimate, or requisite diversity in the organisation based on different ontological conditions or states and (iii) the use of epistemological acts to achieve transitions between those states as a necessary aspect of innovation.

The landscape of management

The matrix shown in Figure 1 identifies four distinct types of management science, two of which currently dominate thinking and practice, and two of which are emerging from the new sciences of complexity. The vertical dimension looks at the nature of our possible understanding of the system, and the horizontal at means of controlling or directing that system. In the vertical dimension design is contrasted with emergence. By design, I mean the ability of a leadership group or consultant to stand outside the system and design the system as a whole based on desired outcomes, gap analysis, etc. In contrast, with emergence, the system cannot be understood or managed as whole, but only through the interactions of the agents (people, processes, technology, government etc.).

In the horizontal dimension, we contrast rules (which could be restated as 'process') that remove ambiguity with heuristics (which could be restated as 'values') that provide direction with a degree of ambiguity that can adapt to different and changing contexts. There is a design element to emergent systems, but as will be described later it is not the system that is designed, but the boundaries and attractors that influence the evolution of such systems.

Design

The dominant approach to management science in the 20th century was based on engineering approaches (rule-based design) in which the leadership of the organisation determined objectives and designed processes to achieve those objectives, together with clear measurements of success and failure. Although this movement can be traced back to Taylor's creation of the term 'Scientific Management', its most recent manifestation is in the Business Process Re-engineering (Hammer & Champy, 1993). In the mid-1980s, based on the work of such figures as Peters, Senge and Nonaka, organisations came to realise that aspects of their experience were non-linear in nature, and the role of leadership was not just about creating and managing process but also about aligning the organisation with clearly expressed objectives, often expressed in mission statements and corporate values. Organisation culture becomes an important issue and the focus on KM and organisational learning starts to gain, and finally seize the high ground from engineering approaches. Most (but not all) popular instantiations of Systems Thinking fall into this area, which shifts from the certainty of rules to the ambiguity of heuristics but does not challenge the concept that the system as a whole can be designed. Another interesting contrast here is between different approaches to strategy. Mintzberg *et al*. (1998) identify 10 schools of strategy. The first three, design, planning and positioning, are most popularly represented by Michael Porter and clearly sit within the Engineering quadrant of the matrix. They were '…tailor made for consultants. They can arrive cold, with no particular knowledge of a business, analyze the data, juggle a set of generic strategies (basic building blocks) on a chart, write a report, drop an invoice and leave'. The remaining six are in various ways related to what I have termed systems thinking, focusing on leadership, learning and environmental factors. Again, one of the more popular representations would be the focus on core competence from Hamel and Prahalad that is closely associated with dominant KM practice both conceptually and temporally.

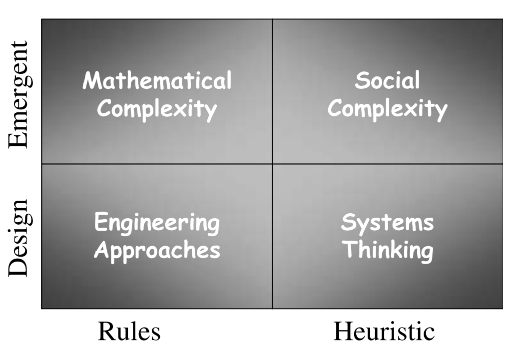

Figure 1 Four types of management science.

Emergent

We are seeing some early signs of strategy formation in the emergent spaces, in particular those taking a cognitive perspective such as Morgan (1986). There have always been indications of common-sense understanding in Mintzberg *et al*'s *Strategy Safari* (1998), although they use the word 'emergence' exclusively in the context of learning systems with no reference to complexity other than indirectly via chaos and evolution. As yet, there is no coherent and certainly no popular formulation of an approach to strategy in the upper quadrants of the matrix.

While it has not yet translated into strategy, we do now have a growing understanding of systems that cannot be designed or understood as a system, but only in respect of the operation of agents within those systems. Here order is emergent and coherence (in terms of causality) retrospective at best, incoherent at worst. Originating from the sciences of complexity and chaos, we are seeing increasing signs of popular acceptance and interest, ranging from sound theory to gross trivialisation. Examples include Axelrod & Cohen (1999), Cilliers (1998), Johnson (2001), Lewin & Regine (1999), Stacy (2001). The existence of such systems based on emergent rather than design-based order allow us to explain much of the failure of current approaches to KM and gain a new insight into new ways of developing methods and tools for the post-SECI period of KM.

The most commonly understood approach to complexity is often associated with the Santa Fe Institute, which, in its popular manifestations (but not exclusively), tends to focus on agent based modelling, referenced in the above matrix as 'mathematical complexity'. Based on work in areas such as insect behaviour, crystal formation and the like, computer-based simulations allow the behaviour of agents to be modelled through mass simulations. This is already being used in aspects of supply chain management, urban design and many other areas. Mathematical complexity shifts from design of the system as a whole to understanding that order emerges from interactions, but it remains with the domain of process and rules. My reason for arguing this is the strong focus on modelling agent behaviour. One of the more readable books addressing this approach in the context of management is *Harnessing Complexity* (Axelrod & Cohen, 1999). The authors outline a clear set of key terms and concepts that we can summarise as follows. The prime unit is an agent (an individual or a social grouping) that can interact with its environment, which includes other agents. Each agent has a strategy that determines what choices that agent will make. Strategies may evolve as they succeed or fail based on some measure of success. Strategies may be copied based on the success of other agents. The necessity of a population of agents is central as it increases the number of potential learning interactions. Variation and interactions lead to adaptation and co-evolution from which order emerges that may then be harnessed. The authors draw on game theory and aspects of computer-based simulation and modelling to create a general theory of management that recognises that 'we all must intervene in Complex Adaptive Systems daily. We all face situations where the classical approach of formulating alternative actions and their likely consequences assumes more understanding and predictive power than we actually have'. They claim that their 'framework shows how the accumulating scientific insights into variation, interaction and selection fit together and can be used to harness complexity'.

Social complexity is the newest entrant to the field, sharing emergence with its mathematical cousin, but arguing that human systems are distinguished from other complex systems. Probably, the best-known approach is participative complexity based on work surrounding Ralph Stacy at the University of Hertfordshire. In the introduction to Griffin's excellent book on leadership (Griffin, 2002), the series editors (Stacy, Griffin and Shaw) identify the dominant voice of management as speaking 'the language of design, regularity and control. In this language, managers stand outside the organizational system, which is thought of as an objective, pre-given reality that can be modelled and designed, and they control it. Managers here are concerned with the functional aspects of a system as they search for causal links that promise sophisticated tools for predicting its behaviour. The dominant voice talks about the individual as autonomous, self-contained, masterful and at the centre of an organisation.' They also identify that much of what I have called the mathematical school of complexity 'talk in a language that is immediately compatible with this dominant voice' seeing complex systems as 'objective realities that scientists can stand outside of and model' furthering the goal of human control. This view is reflected in the matrix by placing mathematical complexity firmly in the rule-based group with engineering. Participative complexity calls a plague on houses of both design schools and that of mathematical complexity arguing that 'humans are themselves members of the complex networks that they form and are drawing attention to the impossibility of standing outside of them in order to objectify and model them. With this intersubjective voice people speak as subjects interacting with others in the co-evolution of a jointly constructed reality. These voices emphasize the radically unpredictable aspects of self-organising processes and their creative potential. These are the voices of decentred agency, which talk about agents and the social world in which they live as mutually created and sustained'.

I have quoted fairly extensively in respect of emergent approaches, partly because they will be less familiar than those based on system level design, partly because there is a tendency in popular thinking to lump chaos and complexity together when they are different in that, to quote Axelrod and Cohen 'chaos deals with situations such as turbulence that rapidly become disordered and unmanageable. On the other hand complexity deals with systems composed of many interacting agents. While complex systems may be hard to predict, they may also have a good deal of structure and permit improvement by thoughtful intervention'. A worse confusion is between systems thinking and complexity, often made simply because both use some of the same language and deal with non-linear causality; the error of this confusion is well handled by Griffin (2002).

In effect, at a simplistic level we could characterise each of the four quadrants of the matrix by four sets of authors, two popular, two less well known, all of whom are referenced above. Moving anti-clockwise from the bottom left they are Hammer, Senge, Stacy and Axelrod. While all four quadrants have various things in common, they are based on different understandings and theories about the nature of the organisation and its environment and appear antithetical to each other, although each appears in context to make sense. That apparent paradox leads us to a proposed *Contingent Complexity*, which will be more fully described later.

The implications for the knowledge-based enterprise

Most current approaches to KM are clearly based on both design approaches and in the common speak of conferences and practitioners, the difference between rules and heuristics is often expressed at that between culture and technology. It is not infrequent to hear quotes such as 'a good KM solution is 10% technology and 90% culture' or variations on that theme. To a large extent, this is a false dichotomy in that human culture is at least in part formed by our capacity to use and create tools and modern technology just happens to be a particularly powerful tool. A good tool fits our hand when we pick it up; too much of the day-to-day practice of KM technology-based solutions requires the user to bio-re-engineer theirs hands, or in extreme cases their brains to fit the tool. Part of this problem stems from the design metaphor itself. In practice, most organisations have accommodated both rule- and heuristic-based approaches by first of all designing a tool based on an idealised model of knowledge creation or exchange and having implemented that tool proceed to create a cultural change programme to align the organisation with the tool. Given that a large part of human knowledge is emergent and not susceptible to design, this produces an inevitable disconnect and results in either abandonment of the formal management of knowledge, or an attempt to enforce 'best practice' more vigorously. The latter is dangerous in that humans tend to accommodate to formal process by 'work arounds' and camouflage behaviour, which gives the impression that a system is working, when in fact it is not (Snowden, 2002). In my more polemical moments, I characterise the technology and cultural approaches to KM as a contest between the *techno- fetishists* for whom humans are an inconvenience in the way of their wonderfully engineered solutions and the *new age fluffy bunnies*, who despise process and think that KM is all about human interaction. In effect, both are right and both are wrong, but neither is complete. To deny the role of technology in KM would be foolish, as would be to deny the complexity of human interaction.

The position that I wish to argue is very simple and runs as follows:

Engineering approaches have become a hygiene factor in KM

The early period of formal KM where the focus was on tacit to explicit knowledge went in parallel with the wide-scale adoption of process re-engineering and adopted many of its forms. The increasing scalability of basic collaboration technologies, originally pioneered by Lotus and others, made it possible to link and connect people in a different way across time and space. The focus on codification was inevitable, and also brought major benefits. It is now something that we know how to do and there is enough experience for conventional academic research to study, interpret and thereby improve the field. Link that with the ubiquitous use of web-based solutions and it would be difficult to argue that KM has not transformed the landscape in both Government and Industry alike. In consequence there is no real excuse for not doing it well. It is about defining and enforcing best practice and now part of the day-to-day fabric of an organisation.

I should also add as a sidebar, that not all engineering or process necessarily requires technology. To take one example the company handbook; older readers may remember this august ring bound volume with its ritual of quarterly amendments, taking out one page, replacing it with another, changing this word, changing that. Now replaced with web-based material in most organisations, the assumption is too frequently made that people will constantly check the web site for updates and changes. In practice, they do not until a mistake or error triggers a need for access. The older manual version had the advantage of ritual in that it created a cyclical awareness of change that instantiated learning in practice.

In summary, engineering approaches do not provide any new or novel insight or innovative capability other than as a hygiene factor, not doing it well results in damage and loss.

Systems thinking starts to explore an interesting space but is limited in application

Much of *The Knowledge Creating Company* (Nonaka & Takeuchi, 1995) is rooted in Systems Thinking and it is ironic that the SECI model of tacit to explicit knowledge transformation was utilised as a rule-based engineering model when it was intended as something less mechanical. Nonaka attempted to rectify some of the abuse with the introduction of the concept of *'ba'* (Nonaka & Konno, 1998), but this did not really disrupt the engineering takeup. I have used systems thinking as a catch all title here, the most popular approach is associated with Senge (1990), who linked both systems thinking with ideas such as single- and double-loop learning to create the approach popularly known as the Learning Organisation (LO).

It is not surprising given the close linkage between knowledge and learning, to see LO concepts and practices becoming closely linked with KM. In particular, the growing impact of Human Resource Departments in KM practice, often competing with Information Technology for ownership of KM, has accelerated this process. One of the key understandings from this approach is the recognition that there are many options in human interactions and that choices have to be made; this could be characterised as good practice, in contrast to the best practice approach common in Engineering. Those choices, normally the function of the leader as a designer (readers may remember Senge's metaphor of the ship) have to be made, and then the culture, or people aspect of the organisation has to be aligned with those objectives. There are limitations to this approach that have been best articulated by the school of participative complexity. However, there is still value to be gained in these approaches for KM. In contrast with emergent approaches, they offer some stability and structure, with the possibility of feedback loops and a degree of management control. There are many aspects of a KM programme for which this is appropriate where social convention and practice limits the range of interactions to the point where a system can be designed. A good example would be the widespread use of matrix-based organisational structures that introduce non-linearity into management, but a way that is constrained to the point where we can understand the system as a system and are not subject to the vagaries of emergent behaviour. Humans are very good at this; we do it at the level of the government as well as in communities and organisations at all levels.

However, as for Engineering the methods and concepts associated with this approach are now reasonably well known and available. Again they are predominantly a hygiene factor; however, they do carry an additional danger in that they often appear to account for behaviour that is emergent and their design base makes them attractive to traditionally trained managers seeking a higher degree of control than is possible with emergent phenomena. While the boundaries of engineering are reasonably clear, the legitimate boundary of systems thinking is more fussy and the language more ambiguous and thus requires greater attention to be paid. How to do this will be elaborated in Part II of this paper.

Mathematical complexity gives us new tools, but is not a management concept

Agent-based modelling and a range of techniques based on genetic algorithms, all the sophistication of natural language processing and complex search engines are readily accepted within the KM world. The sophistication and adoption of these techniques is going to increase over time and the main danger is creating an over-dependency on the tools at the cost of human judgement. Used as a construct for management, mathematical complexity offers little advantage over and above the engineering approaches from which it is a logical development. It can too often confuse simulation with prediction. The fact that I can produce a simulation of bird flocking behaviour through modelling a system in which the agents fly to the centre of the flock, match speed and avoid collision, does not mean that I can predict if the next flock of birds will go to the left or the right of a tree until the moment of its happening, and a small change in initial conditions can give rise to a major change in outcomes. In a sense the dependence on process can destroy human capabilities by breeding in conformity. One of the things that human beings are good at is managing for serendipity or seeing connections between things or 'conceptual blending' (Fauconnier & Turner, 2002). In a very real sense, humans act to create structure in their complex environments, and invent by making conceptual connections that would not be made in nature. This is a part of one of the distinguishing characteristics of humans from other life forms, namely our ability to create 'scaffolding' (Clark, 1997) in the external world to hold our knowledge and create structured, sometimes non-emergent means of organising the world.

Another danger in taking an approach based purely on mathematical complexity is the blind application of Darwinian evolution – culling the bottom 10% of performers and other practices that fail to realise the interactive, participatory aspects of human interaction. Sometimes this manifests itself in a quasi-anarchistic model of self-determination and libertarianism which in a post-Enron should not be considered viable.

Social complexity offers great promise if it is not taken to extreme

Social complexity offers interesting possibilities. Unlike the design-based approaches it is not a hygiene factor, and given the infinite number of possibilities that arise is social complexity is unlikely ever to be so. So the basic argument that will be taken up in the second part of this article is that the main competitive advantage, or for the public sector the greatest opportunity for low cost – high value service provision, lies not in marginal improvements in design, but in opening up the new possibilities that social complexity allows. Interestingly while other approaches focus on best or good practice, social complexity allows us to build sustainable systems based on worst practice; avoiding failure providing a more natural learning mechanism than imitating success under conditions of uncertainty.

The extreme form of social complexity is that of participative complexity. Powerful in its recognition of the importance of conversation and socially constructed meaning, potentially blind to the power and advantages to be gained from engineering and systems thinking-based approaches, within boundaries of legitimacy. Accepting this issue of boundaries is the basis of *contingent complexity* in which complexity in human systems is held to be contingent on context. This approach and its consequences for KM in the context of innovation will be the subject of the second part of this article.

Acknowledgments

This matrix in figure 1 was developed in part for a forthcoming report to the European Commission on the possibilities for Europe to gain strategic advantage as a knowledge economy. The contribution of Max Boisot, Rod Shelton, Jonathan Sage and Peter Stanbridge in many conversations is hereby acknowledged.

References

AXELROD R and COHEN M (1999) *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. The Free Press, New York.

CHRISTENSEN CM and RAYNOR ME (2003) Why hard-nosed executives should care about management theory. *Harvard Business Review* **81(9),** 67–74.

CILLIERS P (1998) *Complexity and Postmodernism*. Routledge, London.

CLARK A (1997) *Being There: Putting Brain, Body, and the World Together Again*. MIT Press, Cambridge, MA.

FAUCONNIER G and TURNER M (2002) *The Way that We Think: Conceptual Blending and the Mind's Hidden Complexities*. Basic Books, New York.

GRIFFIN D (2002) *The Emergence of Leadership*. Routledge, London.

JOHNSON S (2001) *Emergence*. Penguin, London.

HAMMER M and CHAMPY J (1993) *Reengineering the Corporation*. Nicholas Brealey Publishing, London.

KURTZ C and SNOWDEN D (2003) The New Dynamics of Strategy: sense making in a complex-complicated world. *IBM Systems Journal* **42(3),** 462–483.

LEWIN R and REGINE B (1999) *The Soul at Work*. Orion, London.

MINTZBERG H, AHLSTRAND B and LAMPEL J (1998) *Strategy Safari*. Prentice-Hall, Englewood Cliffs, NJ.

MORGAN G (1986) *Images of Organisation*. Sage, London.

NONAKA I and KONNO N (1998) The concept of 'Ba': building a foundation for knowledge creation. *California Management Review* **40(3),** 40–54.

NONAKA I and TAKEUCHI H (1995) *The Knowledge-Creating Company*. Oxford University Press, Oxford and New York.

NONAKA I and TOYAMA R (2003) The knowledge-creating theory revisited: knowledge creation as a synthesizing process. *Knowledge Management Research and Practice* **1(1),** 2–10.

POLANYI M (1962) Terry Lectures at Yale University republished in 1983 as *The Tacit Dimension*. Doubleday, Garden City, NY.

SENGE P (1990) *The Fifth Discipline: The Art and Practice of the Learning Organisation*. Doubleday, New York.

SNOWDEN D (2002) Complex acts of knowing: paradox and descriptive self awareness. *Journal of Knowledge Management* **6(2),** 100–111.

STACY R (2001) *Complex Responsive Processes in Organizations: Learning and Knowledge Creation*. Routledge, London.

About the authors

Dave Snowden is the Director of IBM's Cynefin Centre for Organisational Complexity, based at the University of Cardiff, which focuses on the application of complexity theory to organizational issues. He has an MBA from Middlesex University and a BA in Philosophy from Lancaster University. He is adjunct Professor of Knowledge Management at the University of Canberra, an honorary fellow in knowledge management at the University of Warwick and teaches in MBA programmes at Warwick, Sophia Antipolis and Piacenza. He regularly consults at the board level with some of the world's largest companies as well as to Government and NGOs. In addition, he sits on a number of advisory and other bodies, including the British Standards Institute committee on standards for Knowledge Management.