

The new dynamics of strategy: Sense-making in a complex and complicated world

by C. F. Kurtz
D. J. Snowden

In this paper, we challenge the universality of three basic assumptions prevalent in organizational decision support and strategy: assumptions of order, of rational choice, and of intent. We describe the Cynefin framework, a sense-making device we have developed to help people make sense of the complexities made visible by the relaxation of these assumptions. The Cynefin framework is derived from several years of action research into the use of narrative and complexity theory in organizational knowledge exchange, decision-making, strategy, and policy-making. The framework is explained, its conceptual underpinnings are outlined, and its use in group sense-making and discourse is described. Finally, the consequences of relaxing the three basic assumptions, using the Cynefin framework as a mechanism, are considered.

Over the past several years, our group has been conducting a program of disruptive action research using the methods of narrative and complexity theory to address critical business issues.¹ Action research has been defined as grounding theory in contextual exploration, emphasizing participation, and embracing change.

We started work in the areas of knowledge management, cultural change and community dynamics, then expanded into product development, market creation and branding, and in recent years have been working increasingly in the area of national and organizational strategy. Some of this work has been

directly funded by the U.S. government through DARPA (Defense Advanced Research Project Agency) as well as by other government agencies (in particular in Singapore) which are interested in new approaches to supporting policy-making. The central element of our approach is the Cynefin framework for sense-making. In this paper we describe the framework and its conceptual basis, and we detail some of its uses for sense-making to support decision-making in varied dynamical contexts.

Conceptual approach. We begin by questioning the universality of three basic assumptions that pervade the practice and to a lesser degree the theory of decision-making and policy formulation in organizations. These are:

The assumption of order: that there are underlying relationships between cause and effect in human interactions and markets, which are capable of discovery and empirical verification. In consequence, it is possible to produce prescriptive and predictive models and design interventions that allow us to achieve goals. This implies that an understanding of the causal links in past behavior allows us to define “best practice” for future behavior. It also implies that there must be a right or ideal way of doing things.

The assumption of rational choice: that faced with a choice between one or more alternatives, human ac-

©Copyright 2003 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

tors will make a “rational” decision based only on minimizing pain or maximizing pleasure; and, in consequence, their individual and collective behavior can be managed by manipulation of pain or pleasure outcomes and through education to make those consequences evident.

The assumption of intentional capability: that the acquisition of capability indicates an intention to use that capability, and that actions from competitors, populations, nation states, communities, or whatever collective identity is under consideration are the result of intentional behavior. In effect, we assume that every “blink” we see is a “wink,” and act accordingly. We accept that we do things by accident, but assume that others do things deliberately.

This paper contends that although these assumptions are true within some contexts, they are not universally true. We also believe that in decision-making at both policy-making and operational levels, we are increasingly coming to deal with situations where these assumptions are not true, but the tools and techniques which are commonly available assume that they are.

Order and chaos in antiquity. The human distinction between order and chaos goes back to an abundant presence in mythology, in which order arises out of (and thus requires) and then vanquishes (and thus destroys) the mysterious forces of chaos. For example, in the *Enuma Elish*,² the Babylonian epic of creation, the world began under the reign of Tiamat, the mother of all things. In Tiamat’s world, “none bore a name, and no destinies were ordained.” After several generations, Tiamat’s god-children appointed a champion to seize control. Marduk not only defeated his ancestor, but “split her up like a flat fish into two halves” that became heaven and earth. He then proceeded to order the universe in finer and finer detail:

He [Marduk] made the stations for the great gods;
The stars, their images, as the stars of the Zodiac,
he fixed.
He ordained the year and into sections he divided
it;
For the twelve months he fixed three stars. . . .
He founded the station of Nibir [the planet Jupiter]
to determine their bounds;
That none might err or go astray . . .

Note the words “fixed,” “ordained,” “divided,” “determined,” “err,” and “astray.” Control (in the first

four terms) and an absolute knowledge of right and wrong (in the last two) are the salient points of Marduk’s new world. Of course Tiamat was never entirely vanquished; forces of chaos appear in all traditions in the form of tricksters and malcontents such as Bacchus, Loki, Coyote, the Monkey King, Anansi, and Hermes. The forces of order and chaos danced with each other throughout ancient times.

Science, order, and epiphenomena. Aristotle defined four types of cause: the material (what you are made of, your muscles and organs), the efficient (how you came to be, the fact that your parents gave birth to you), the formal (your type, your species), and the final (your function, your life itself, your place in the universe). He believed that to understand an event or entity, one had to consider all of these factors in the particular, the mysterious as well as the ordered. As others have pointed out,³ the focus of Western thinking post Kant on efficient cause only is to the detriment of knowledge.

Since the birth of enlightenment science, the distinction between order and chaos has held a prominence that has profoundly influenced conceptual and practical thinking. Kant separated things that we can know empirically from things that are the province of God, and thereby helped to section off all but efficient causes to epiphenomena that could be safely ignored. This concept of ordered science triggered a massive growth in human knowledge and extended over many disciplines. For example, sociology grew out of philosophy partly in an attempt to create a “science of society” that could duplicate the advances being made in physics and biology through systematic observation and causal explanation. It was argued, mainly by Comte, that it was theoretically possible to discover laws similar to those of physics which could explain the behavior of people in societies. The growth of technology and the dominance of engineering-based approaches arising from the need for automation and scalability reenforced the desire for and the assumption of order. In popular literature, the belief that all things can be known (in a Newtonian sense) persisted well into this past century. Asimov’s classic science fiction *Foundation Trilogy*⁴ builds on the character of Hari Seldon, the founder of psychohistory, whose mathematics permits the prediction of human behavior and social change centuries into the future.

The development of management science, from stopwatch-carrying Taylorists⁵ to business process reengineering, was rooted in the belief that systems were

ordered; it was just a matter of time and resources before the relationships between cause and effect could be discovered.⁶ The case study approach of many M.B.A. programs and the desire for precise recommendations from policy teams and external consultants perpetuate the underlying assumption of universal order. Good leadership is linked to certain competences that (it is claimed) can be mapped and identified, and then replicated. The desire for order can even lead people to accept completely abhorrent working conditions and political structures simply to avoid “chaos” (a tendency exploited by dictators from Pisistratus to Hitler and beyond). All of these approaches and perceptions do not accept that there are situations in which the lack of order is not a matter of poor investigation, inadequate resources, or lack of understanding, but is a priori the case—and not necessarily a bad thing, either.

Complexity science. A new awareness of the ancient counterpart to order began over a century ago with Poincaré and several others, and has surged in recent decades.⁷⁻¹⁰ In fact there is a fascinating kind of order in which no director or designer is in control but which emerges through the interaction of many entities. Emergent order has been found in many natural phenomena: bird-flocking behavior can be simulated on a computer through three simple rules¹¹; termites produce elegant nests through the operation of simple behaviors triggered by chemical traces¹²; each snowflake is a unique pattern arising from the interactions of water particles during freezing.¹³ The patterns that form are not controlled by a directing intelligence; they are self-organizing.¹⁴ The new science of complexity spawned by these findings is interdisciplinary, touching fields from mathematics to evolution to economics to meteorology to telecommunications.¹⁵ In the domain of emergent order, the goal “to predict (and thereby control) the behavior of systems not yet studied (but similar to those that have been studied) under conditions not yet extant and in time periods not yet experienced”¹⁶ is difficult if not impossible to achieve—but other goals are achievable.

Awareness of emergent order has as yet had comparatively little influence on mainstream theory and practice in management and strategy, though some authors have been exploring these topics with some success.¹⁷⁻²¹ We are sorry to say that in our opinion too many of the books written for a popular business audience on the subject have been marred by misunderstandings, misapplications, and most of all misplaced zeal (some even falling into pseudo-sci-

entific and pseudo-religious thinking). Our group has been using ideas based on complexity science in our action research work for the past several years, and the work we will describe here benefits from that influence.

Contextual complexity. A considerable amount of research and some early practice is taking place using complex system principles, mainly using computing power to simulate natural phenomena through agent-based models.¹² Well-known examples include routing optimization of airfreight and telecommunication signals based on large volume data modeling in which each agent is programmed to operate on simple rules, the result of which is the emergence of complex patterns of behavior.²² We believe that such tools are valuable in certain contexts, but are of more limited applicability when it comes to managing people and knowledge. There are at least three important contextual differences between human organizations and those of ant colonies that make it more difficult to simulate them using computer models. In a sense, because we have not seen these issues addressed as fully as we would like, we put them forth as challenges to agent-based simulation of human behavior.

Humans are not limited to one identity. In a human complex system, an agent is anything that has identity, and we constantly flex our identities both individually and collectively. Individually, we can be a parent, sibling, spouse, or child and will behave differently depending on the context. Collectively, we might, for example, be part of a dissenting community, but in the face of a common threat, we might assume the identity of the wider group. Accordingly, it is not always possible to know which unit of analysis we are working with.

There are generally three solutions to this problem of the unit of analysis in the social simulation literature.²³ First, individuals are modeled and group behavior is explained by the concept of “norms.”²⁴ But identity goes deeper than norms—it determines not only reactions but perceptions and patterns of experience. A second solution is to model groups as agents, but much internal diversity and patterning is suppressed in this approach. It also makes assumptions about intent and predictability that are difficult to sustain. The third solution to the unit-of-analysis problem is to consider the rule or idea or “meme” as the unit of analysis,²⁵ but again we find that insufficient to capture the dynamics of multiple identities. We would like (but do not expect) to see sim-

ulations of human behavior able to encompass multiple dynamic individual and collective identities acting simultaneously and representing all aspects of perception, decision-making, and action.

Humans are not limited to acting in accordance with predetermined rules. We are able to impose structure on our interactions (or disrupt it) as a result of collective agreement or individual acts of free will.

**We are capable of shifting
a system from complexity
to order and maintaining it there
in such a way that
it becomes predictable.**

We are capable of shifting a system from complexity to order and maintaining it there in such a way that it becomes predictable. As a result, questions of intentionality play a large role in human patterns of complexity.³ It is difficult to simulate true free will and complex intentionality (for example, retrospective elaboration, duplicity, groupthink, rumor, self-deception, manipulation, surprise, confusion, internal conflict, stress, changes in the meanings of previously unambiguous messages, the deliberate creation of ambiguity, inadvertent disclosure, charisma, cults, and pathologies) within a rule-based simulation. Social simulations have addressed issues such as cooperation, reputation, gossip, lying, and trust,²⁶ but always within an artificial framework which allows only limited numbers of options and considers limited numbers of phenomena operating at once. It is interesting that searching the Internet for “simulation” and most of the terms listed above brings up instances of simulations with which a user interacts to explore patterns.²⁷ This may represent a gap between agent-based simulation and human behavior similar to that found by the “strong artificial intelligence (AI)” school of the 1970’s whose goal was to reproduce human intelligence.²⁸ We do not mean to say that there is no value to simulation of human behavior, but we do think we should not expect it to succeed any time soon in predicting what people will do in any particular circumstance.

Eventually the concept of “intelligent augmentation”^{29–30} grew in popularity and could be said to bridge the gap between strong AI and reality. It is possible that a middle ground between the belief that

all human behavior can be simulated and the belief that the very effort is questionable will be found in the use of simulation not to explain or imitate but to support human decision-making.

Humans are not limited to acting on local patterns. People have a high capacity for awareness of large-scale patterns because of their ability to communicate abstract concepts through language, and, more recently, because of the social and technological infrastructure that enables them to respond immediately to events half a world away. This means that to simulate human interaction, all scales of awareness must be considered simultaneously rather than choosing one circle of influence for each agent. There is also the matter of simulating the interaction (conflict, reinforcement) between local and global awareness. Many of the emergent patterns we see in nature depend critically on the limited (that is, local) ability of activators to diffuse through a viscous medium.¹³ We have not yet seen addressed how these issues cause complex patterns in human societies to differ from complex patterns in systems of locally aware agents.

We call our practice of keeping the human context foremost in our considerations phenomenological or more commonly “contextual complexity.” It means mainly that when we use agent-based simulation (and we do, in certain circumstances), we use it as a tool for the exploration of possibility and generation of ideas, not as a tool for recommending courses of action.

Order and un-order. To avoid much repetition of the longer terms “directed order” and “emergent order,” we call emergent order “un-order.” Un-order is not the lack of order, but a different kind of order, one not often considered but just as legitimate in its own way. Here we deliberately use the prefix “un-” not in its standard sense as “opposite of” but in the less common sense of conveying a paradox, connoting two things that are different but in another sense the same. Bram Stoker used this meaning to great effect in 1897 with the word “undead,” which means neither dead nor alive but something similar to both and different from both. According to R. D. Cureton, e. e. cummings also used the prefix this way in his poetry. Says Cureton,³¹ “In normal usage, being and existing are stative concepts. They are not actions which a person must consciously perform, engage in, create. Words such as *unbe* and *unexist*, however, force the reader to see the dynamic nature of human existence. . . .” Thus by our use of the term

“un-order,” we challenge the assumption that any order not directed or designed is invalid or unimportant.

Keep the baby, lose the bathwater. Let us sound a quick warning about running into the trap of believing that everything is complex. Some recent popular books on complexity in business and management have been full of breathy enthusiasm for the “edge of chaos” and would have businesses maintain themselves as far from equilibrium as possible, regardless of context or purpose. We think this is throwing out the baby with the bathwater. We cannot simply go from saying “things are ordered” to saying “things are un-ordered” and leave it at that; things are both ordered and un-ordered at once, because in reality order and un-order intertwine and interact. Kostof³² puts it well in his description of cities: “. . . the two primary versions of urban arrangement, the planned and the ‘organic,’ often exist side by side. . . . Most historic towns, and virtually all those of metropolitan size, are puzzles of premeditated and spontaneous segments, variously interlocked or juxtaposed. . . .” In other words, it is useful to artificially separate order and un-order so that we can understand the different dynamics involved, but we should not expect to find one without the other in real life. In many organizations, for example, formal command structures and informal trust networks support (while simultaneously competing with) each other.³³⁻³⁴ The joke that “the only thing worse than an inefficient bureaucracy is an efficient bureaucracy” has some ground in reality.

Methods for un-ordered space

Ordered-systems thinking assumes that through the study of physical conditions, we can derive or discover general rules or hypotheses that can be empirically verified and that create a body of reliable knowledge, which can then be developed and expanded. As we have mentioned, this assumption does not hold in the domain of un-order.

In practice, all decision makers know this: however much they might like things to be ordered, they know that there are also circumstances in which “cultural factors,” “inspired leadership,” “gut feel,” and other complex factors are dominant. All of these are patterns, which arise through the interaction of various entities through space and time. In the space of un-order the seeds of such patterns can be perceived, and new ways of thinking can emerge. In fact, learning to recognize and appreciate the domain of un-

order is liberating, because we can stop applying methods designed for order and instead focus on legitimate methods that work well in un-ordered situations. Tom Stewart³⁵ references the case of a group of marines taken to the New York Mercantile Exchange in 1995 to be taught and to play with simulators of the trading environment. Naturally the traders won each time. But when the traders visited the Marine Corp’s base in Quantico and played war games against the marines, they won yet again. What they realized is that the traders were skilled at spotting patterns and intervening to structure those patterns in their favor. The Marines, on the other hand, like most business school graduates, had been trained to collect and analyze data and then make rational decisions. In a dynamic and constantly changing environment, it is possible to pattern un-order but not to assume order.

In another case, a group of West Point graduates were asked to manage the playtime of a kindergarten as a final year assignment. The cruel thing is that they were given time to prepare. They planned; they rationally identified objectives; they determined backup and response plans. They then tried to “order” children’s play based on rational design principles, and, in consequence, achieved chaos. They then observed what teachers do. Experienced teachers allow a degree of freedom at the start of the session, then intervene to stabilize desirable patterns and destabilize undesirable ones; and, when they are very clever, they seed the space so that the patterns they want are more likely to emerge.

In the ordered domain we focus on efficiency because the nature of systems is such that they are amenable to reductionist approaches to problem solving; the whole is the sum of the parts, and we achieve optimization of the system by optimization of the parts. In the domain of un-order, the whole is never the sum of the parts; every intervention is also a diagnostic, and every diagnostic an intervention; any act changes the nature of the system. As a result, we have to allow a degree of sub-optimal behavior of each of the components if the whole is to be optimized.

Pattern entrainment

Humans use patterns to order the world and make sense of things in complex situations. Give a child a pile of blocks, and he or she will build patterns out of them. Give an adult a daily commute, and he or she will build patterns within it. Patterns are some-

thing we actively, not passively, create, as Mary Douglas³⁶ so well pointed out:

. . . whatever we perceive is organized into patterns for which we the perceivers are largely responsible. . . . As perceivers we select from all the stimuli falling on our senses only those which interest us, and our interests are governed by a pattern-making tendency, sometimes called a schema. In a chaos of shifting impressions, each of us constructs a stable world in which objects have recognizable shapes, are located in depth and have permanence. . . . As time goes on and experience builds up, we make greater investment in our systems of labels. So a conservative bias is built in. It gives us confidence.

Visually, we hold in sharp focus at any one instant a mere tenth of a percent of our visual range, so even the process of seeing is one of putting together many disparate observations.³⁷ We fill in the gaps to create an experience-based pattern on which we act. This aspect of human decision-making is a great source of power, but it also brings limitation.

A television advertisement for a liberal broadsheet newspaper in the United Kingdom illustrates this well. The advertisement is set in a terraced inner city street. It is dusk; litter blows down the street; overall, a threatening environment.

Scene One: The camera points down the street and picks up a skinhead who comes around the corner. A police car stops beside the skinhead, who immediately runs towards the camera. Pattern entrainment of decision-making: what do we assume? The skinhead must be running from the police.

Scene Two: The camera now changes its perspective; it is now behind the skinhead, and we see that he is running towards a well dressed man who is clutching a briefcase—which must be full of money, as the man conforms to the stereotype of a rent collector. We can also see that the man is terrified in the face of the rapidly approaching skinhead. Pattern entrainment of decision making: what do we assume? The man is going to be mugged.

Scene Three: The camera changes perspective for the final time; we are now looking down on the street from above, and we see the skinhead grab

the man and pull him into the portico of a building just before a crate of building material would have fallen on his head and killed him.

The unspoken message of a brilliant advertisement is, “See things from a different perspective; read the newspaper.”

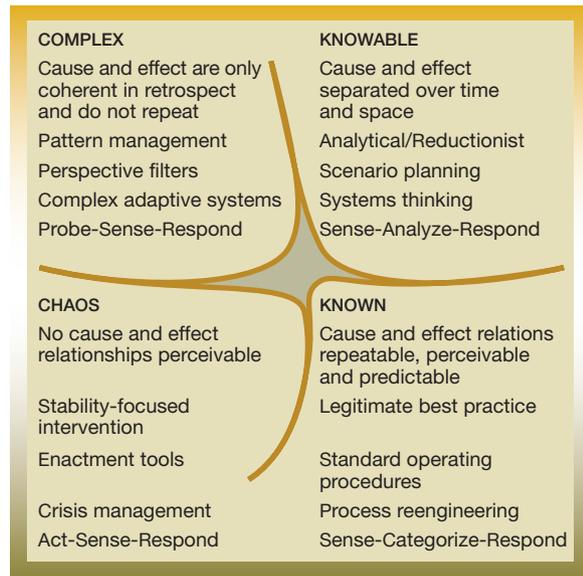
Most of the time if you are standing in a run-down inner city area at dusk carrying a briefcase full of money, and a skinhead suddenly runs towards you, it is not a good idea to stand there and say “Ah, I may be about to be rescued from a crate of falling building material”; you should run for your life. The issue in decision-making is to know when to run and when to stand still. A choice must be made between allowing the entrained patterns of past experience to facilitate fast and effective pattern application and gaining a new perspective because the old patterns may no longer apply. We will address these differences in the following section.

The Cynefin framework

The name *Cynefin* is a Welsh word whose literal translation into English as habitat or place fails to do it justice. It is more properly understood as the place of our multiple affiliations, the sense that we all, individually and collectively, have many roots, cultural, religious, geographic, tribal, and so forth. We can never be fully aware of the nature of those affiliations, but they profoundly influence what we are. The name seeks to remind us that all human interactions are strongly influenced and frequently determined by the patterns of our multiple experiences, both through the direct influence of personal experience and through collective experience expressed as stories.

The Cynefin framework originated in the practice of knowledge management as a means of distinguishing between formal and informal communities, and as a means of talking about the interaction of both with structured processes and uncertain conditions. It has now outgrown its application in knowledge management, having been in use by our group for several years in consultancy and action research in knowledge management, strategy, management, training, cultural change, policy-making, product development, market creation, and branding. We are now beginning to apply it to the areas of leadership, customer relationship management, and supply chain management, with other topics to come. It has also been used by third parties.³⁸

Figure 1 Cynefin domains



We consider Cynefin a *sense-making* framework, which means that its value is not so much in logical arguments or empirical verifications as in its effect on the sense-making and decision-making capabilities of those who use it. We have found that it gives decision makers powerful new constructs that they can use to make sense of a wide range of unspecified problems. It also helps people to break out of old ways of thinking and to consider intractable problems in new ways. The framework is particularly useful in collective sense-making, in that it is designed to allow shared understandings to emerge through the multiple discourses of the decision-making group.

We make a strong distinction here between sense-making frameworks and *categorization* frameworks. In a categorization framework, four quadrants are often presented in a two-by-two matrix (for examples, pick up any management textbook or analyst report). Typically, it is clear (though often unstated) that the most desirable situation is to be found in the upper right-hand quadrant, so the real value of such a framework is to figure out how to get to the upper right. In contrast, none of the domains we will describe here is more desirable than any other; there are no implied value axes. Instead, the framework is used primarily to consider the dynamics of situations, decisions, perspectives, conflicts, and changes in order to come to a consensus for decision-making under uncertainty.

As can be seen in Figure 1, the Cynefin framework has five domains, four of which are named, and a fifth central area, which is the domain of disorder. The right-hand domains are those of order, and the left-hand domains those of un-order.

Ordered domain: Known causes and effects. Here, cause and effect relationships are generally linear, empirical in nature, and not open to dispute. Repeatability allows for predictive models to be created, and the objectivity is such that any reasonable person would accept the constraints of best practice. This is the domain of process reengineering, in which knowledge is captured and embedded in structured processes to ensure consistency. The focus is on efficiency. Single-point forecasting, field manuals, and operational procedures are legitimate and effective practices in this domain. Our decision model here is to sense incoming data, *categorize* that data, and then respond in accordance with predetermined practice. Structured techniques are not only desirable but mandatory in this space.

Ordered domain: Knowable causes and effects. While stable cause and effect relationships exist in this domain, they may not be fully known, or they may be known only by a limited group of people. In general, relationships are separated over time and space in chains that are difficult to fully understand. Everything in this domain is capable of movement to the known domain. The only issue is whether we can afford the time and resources to move from the knowable to the known; in general, we cannot and instead rely on expert opinion, which in turn creates a key dependency on trust between expert advisor and decision maker. This is the domain of systems thinking, the learning organization, and the adaptive enterprise, all of which are too often confused with complexity theory.¹⁸ In the knowable domain, experiment, expert opinion, fact-finding, and scenario-planning are appropriate. This is the domain of methodology, which seeks to identify cause-effect relationships through the study of properties which appear to be associated with qualities. For systems in which the patterns are relatively stable, this is both legitimate and desirable.

Our decision model here is to sense incoming data, *analyze* that data, and then respond in accordance with expert advice or interpretation of that analysis. Structured techniques are desirable, but assumptions must be open to examination and challenge. This is the domain in which entrained patterns are at their most dangerous, as a simple error in an assumption

can lead to a false conclusion that is difficult to isolate and may not be seen.

It is important to note here that by known and knowable we do not refer to the knowledge of individuals. Rather, we refer to things that are known to society or the organization, whichever collective identity is of interest at the time. If I look up my organization's policy on travel to Iceland, I may not know what I will find there, or even how I will find it, but I know I can find something, because I know it is known to the organization. If I want to know how fish sleep, I may not know how or where to find out, but I have a hunch that somebody knows.

Un-ordered domain: Complex relationships. This is the domain of complexity theory, which studies how patterns emerge through the interaction of many agents. There are cause and effect relationships between the agents, but both the number of agents and the number of relationships defy categorization or analytic techniques. Emergent patterns can be perceived but not predicted; we call this phenomenon *retrospective coherence*. In this space, structured methods that seize upon such retrospectively coherent patterns and codify them into procedures will confront only new and different patterns for which they are ill prepared. Once a pattern has stabilized, its path appears logical, but it is only one of many that could have stabilized, each of which also would have appeared logical in retrospect. Patterns may indeed repeat for a time in this space, but we cannot be sure that they will continue to repeat, because the underlying sources of the patterns are not open to inspection (and observation of the system may itself disrupt the patterns). Thus, relying on expert opinions based on historically stable patterns of meaning will insufficiently prepare us to recognize and act upon unexpected patterns.

The decision model in this space is to create *probes* to make the patterns or potential patterns more visible before we take any action. We can then sense those patterns and respond by stabilizing those patterns that we find desirable, by destabilizing those we do not want, and by seeding the space so that patterns we want are more likely to emerge. Understanding this space requires us to gain *multiple perspectives* on the nature of the system. This is the time to "stand still" (but pay attention) and gain new perspective on the situation rather than "run for your life," relying on the entrained patterns of past experience to determine our response. The methods, tools, and techniques of the known and knowable

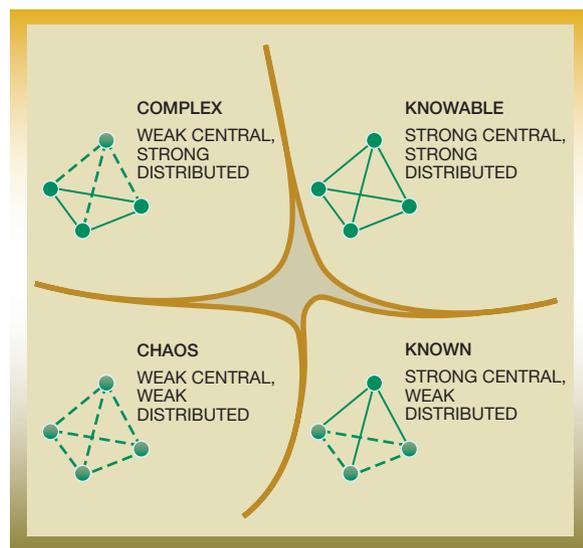
domains do not work here. Narrative techniques are particularly powerful in this space. We have described elsewhere a range of methods designed to stimulate emergent patterns in complex knowledge interactions by increasing the number of perspectives available to a decision maker.³⁹

Un-ordered domain: Chaos. In the first three domains we have described, there are visible relationships between cause and effect. In the chaotic domain there are no such perceivable relations, and the system is turbulent; we do not have the response time to investigate change.⁸ Applying best practice is probably what precipitated chaos in the first place; there is nothing to analyze; and waiting for patterns to emerge is a waste of time. The chaotic domain is in a very real sense *uncanny*, in that there is a potential for order but few can see it—or if they can, they rarely do unless they have the courage to act. In known space it pays to be *canny*, that is, to know how to work the system in all its intricacies (canny meaning not only shrewd but safe). But in chaotic space, a canny ability gets you nowhere (there is no system to be worked). You need a different type of ability, one that is uncannily mysterious, sometimes even to its owner. Canny people tend to succeed in their own lifetimes; uncanny people tend to be recognized and appreciated only centuries later, because during their time their actions appeared to be either insane or pointless. Each of these styles has a unique ability to succeed in a particular space, and each is necessary.

The decision model in this space is to act, quickly and decisively, to reduce the turbulence; and then to sense immediately the reaction to that intervention so that we can respond accordingly. The trajectory of our intervention will differ according to the nature of the space. We may use an authoritarian intervention to control the space and make it knowable or known; or we may need to focus on multiple interventions to create new patterns and thereby move the situation into the complex space. Chaos is also a space we can enter into consciously, to open up new possibilities and to create the conditions for innovation.

The domain of disorder. The central domain of disorder is critical to understanding conflict among decision makers looking at the same situation from different points of view. Often in a group using the Cynefin framework, people agree on what the extremes of the four domains mean in the context they are considering, but disagree on more subtle differ-

Figure 2 Connection strength of Cynefin domains



ences near the center of the space. As a result, individuals compete to interpret the central space on the basis of their preference for action. Those most comfortable with stable order seek to create or enforce rules; experts seek to conduct research and accumulate data; politicians seek to increase the number and range of their contacts; and finally, the dictators, eager to take advantage of a chaotic situation, seek absolute control. The stronger the importance of the issue, the more people seem to pull it towards the domain where they feel most empowered by their individual capabilities and perspectives. We have found that the reduction in size of the domain of disorder as a consensual act of collaboration among decision makers is a significant step toward the achievement of consensus as to the nature of the situation and the most appropriate response.

The apple-orange problem

People are often confused by the apple-orange nature of the four Cynefin domains: they say, “Why not known, knowable, somewhat knowable and unknowable?” or, “Why not simple, complicated, complex and chaotic?” The distinction is intentional. The Cynefin framework is a phenomenological framework, meaning that what we care most about is how people perceive and make sense of situations in order to make decisions; perception and sense-making

are fundamentally different in order versus un-order. The framework actually has two large domains, each with two smaller domains inside. In the right-side domain of order, the most important boundary for sense-making is that between what we can use immediately (what is known) and what we need to spend time and energy finding out about (what is knowable). In the left-side domain of un-order, distinctions of knowability are less important than distinctions of interaction; that is, distinctions between what we can pattern (what is complex) and what we need to stabilize in order for patterns to emerge (what is chaotic). Thus we often draw the framework with the vertical boundaries strong and the horizontal boundaries weak, denoting their relative importance in sense-making.

The Cynefin framework is based on three ontological states (namely, order, complexity and chaos) and a variety of epistemological options in all three of those states. We are currently engaged in further conceptual and experimental work to more strongly develop the separation of ontological from epistemological aspects of the framework in order to root the framework in a variety of scientific disciplines while maintaining the essential interweaving of ontology and epistemology, which appears to be an essential aspect of human sense-making in practice.

Connection strengths of Cynefin domains

Another way to look at the Cynefin framework is in the types of component connections that are most prevalent in each domain (Figure 2). On the side of order, connections between a central director and its constituents are strong, often in the form of structures that restrict behavior in some way—for example, procedures, forms, blueprints, expectations, or pheromones. On the side of un-order, central connections are weak, and attempts at control through structure often fail from lack of grasp or visibility. In the complex and knowable domains, connections among constituent components are strong, and stable group patterns can emerge and resist change through repeated interaction, as with chemical messages, acquaintanceship, mutual goals and experiences. The known and chaotic domains share the characteristic that connections among constituent components are weak, and emergent patterns do not form on their own.

In any of these domains, a reasonable strategy capitalizes on the stability afforded by strong connections without allowing them to harden so much that

they destroy flexibility and also capitalizes on the freedom and renewal afforded by weak connections without allowing them to permanently remove useful patterns.

How the Cynefin framework is used

In our engagements with clients, we use the Cynefin framework at several different levels, depending on the context and purpose of the project. Several complementary exercises and programs help people to internalize and make use of the framework. Some groups consider only the five domains and what sorts of situations or problems can be found there; some consider distinctions between extreme and equivocal conditions within domains; some consider multiple perspectives and how they can be resolved or used to advantage; some talk about boundary transitions, boundary sensing, and boundary management; some talk about dynamics ranging over the whole space of the framework. For example, groups might use the Cynefin framework to gain new insights on a contentious issue, plan interventions to move a situation from one domain to another, consider how they should approach or manage different formal and informal communities, or differentiate their strategies for knowledge retention based on multiple contexts of knowledge exchange.

Contextualization. This is a critical exercise that enters into nearly every Cynefin project. Contextualization is also a good example of how the Cynefin framework concentrates on collective sense-making as a consequence of discourse. We will describe a typical contextualization session here. The session begins with the collection of many items through structured brainstorming. The items might be communities, products, actions, motivations, forces, events, points of view, beliefs, traditions, rituals, books, metaphors, anecdotes, myths, and so on: they are any items that are important to the sense-making process. The items are related to one theme or issue of concern, which should be broad but not infinitely so. Whatever sorts of items are chosen (and multiple types are permissible), they should be diverse and concrete in nature: diverse to allow multiple perspectives to emerge, and concrete to move away from existing entrained abstract beliefs. We assure diversity by giving different groups of people different directions, by giving directions that are deliberately ambiguous and so can be taken in diverse ways, and by changing group compositions frequently so that people do not fall into entrained thinking. To keep items concrete, we rely heavily on narrative

methods. These provide a rich context that allows patterns of experience rather than opinion or belief to emerge.

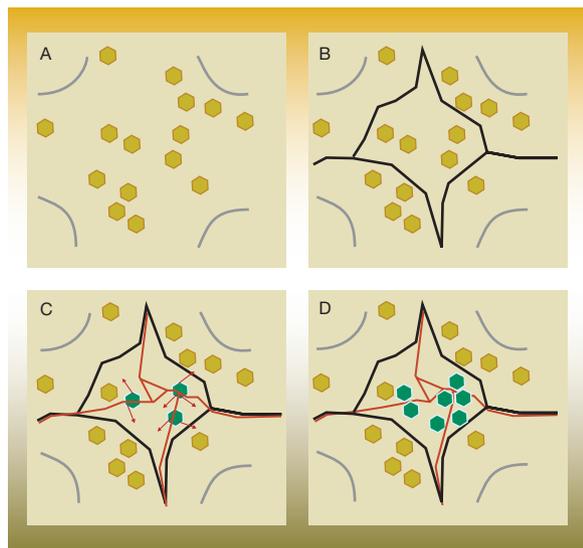
We use an array of different methods to help people prepare a contextually meaningful field of items for sense-making, including:

Narrative database. We might prepare in advance a narrative database of relevant stories surrounding the issue, drawn from oral histories, collected anecdotes, published reports, historical documents, and the like.⁴⁰ Workshop participants might be asked to review the material in advance and answer some sense-making questions about it. In the first part of the workshop, we may ask people to talk about their thoughts on the material and the narratives which they found most compelling. From these discussions a list of such items as situations, actors, events, and forces might be generated.

Convergence methods. Rather than pulling items directly from anecdotes discussed, we might go through one of several convergence methods in order to integrate much disparate material and achieve a stronger set of sense-making items. For example, we might ask people to construct composite fables from anecdotes by using one of a variety of fable templates, working from either the narrative database or their own experiences as source material.⁴¹ This exercise is especially useful when we need to bring people into the realm of fiction so that they can more freely express their true opinions on sensitive subjects. It is also a useful integrator of diverse sources. In a room of 50 people constructing fables in small groups, several hundred anecdotes may be considered as material for creation of each group's fable. Items are then drawn from aspects of the integrated fables (characters, events, situations) for sense-making.

Alternative history. Another useful method (with or without preparatory material) is the use of alternative histories. In this method, we ask people to describe the history of an organization, society, or event, always working backwards from the present (to any starting time they think is appropriate). We ask them to determine turning points (moments when small events caused large changes), write them simply on hexagonal sticky notes, and array them on a wall. (We prefer hexagonal notes because people tend to cluster them in beehive shapes, while people tend to categorize square notes.) When the factual history is complete, we then ask people to choose two extreme

Figure 3 Cynefin sample contextualization session



states, one positive and one negative; each of these must be utterly inconceivable. They then work backwards (always backwards, to avoid entrained “what usually comes next” thinking) with fictional timelines to reach points on the factual timeline. These connection points are not determined in advance, but emerge from the growing reverse fiction. The turning points may then be considered and described, listing, for example, actors, communities, and factors in play during those moments. These become the sense-making items.

Next, as shown in Figure 3A, the group comes together with their assembled sense-making items, perhaps derived from multiple processes, and discusses the extremes (not the domains) of the Cynefin space. They consider the corner where everyone knows the right answer, the corner where an expert could be expected to know the right answer, the corner where the situation only becomes clear retrospectively, and the corner in which there is no right answer. If possible, items are selected from the set that exemplify that dynamical extreme in the context of the issue or events being considered. This is an important step, because it begins to build the framework. We say “build the framework” because the Cynefin framework is created anew each time it is used, with distinctions meaningful to the current context. To some extent, it does not even exist in the way we describe it here, devoid of context, but is always used to en-

able sense-making in a particular setting. (It can also be used to negotiate meaning between contexts, but that use is beyond the scope of this paper.)

After the extreme situations are considered, all of the sense-making items are placed within the overall Cynefin space—without boundaries—where the items seem to fall, using dynamic placement. (We sometimes ask people to think of it as though the item they are placing has four elastic bands connecting it to the four corners, and to find the place where it comes to rest depending on how hard the different bands pull on the item. We then ask them to check the placing of each item in relation to others.) Clusters are allowed to form where items seem similar in the space, but are not required. This is often done in person using large walls and hexagonal sticky notes; it can be done electronically, but must be done socially.

It is important to mention at this point that discussion is encouraged during the placement of items. This is unlike the methods of affinity diagrams, in which people are asked to refrain from speaking.⁴² In our experience, although the proponents of affinity diagrams are correct that silence equalizes the verbal contributions of those who speak up and those who hold back, the placement of items is not equalized by silence: often people just watch very carefully what the boss is doing. In a sense this only drives power differences underground rather than removing them. We find it is better to remove large power differences before contextualization; in fact, we have found that comparing the patterns made by people at different levels of management working on separate contextualizations can be of great value, especially for tackling difficult management issues. Another significant element of the placement of items is that, as in the group construction of affinity diagrams, we encourage people to consider all items together and to keep moving any items they like until the patterns they have produced make sense.

When all the items are placed, as shown in Figure 3B, lines are drawn between hexagons that are clearly in one domain or another, leaving a possibly large central area of disorder. Then, as shown in Figures 3C and 3D, the boundary lines are “pulled in” to make the distinctions between domains more complete. Hexagons that lie on the new lines are bifurcated, trifurcated or quartered, forming new clusters of hexagons on either side of the boundaries. This stage involves much discussion, as consensus has to be reached.

At this point, the characteristics of each domain and boundary—as they pertain to the context at hand—can be considered using the sense-making items in place.

Use of the contextualized framework. After the Cynefin framework has been created, we may move on to other exercises in which the items in specific areas (for example, near the boundaries or at the extremes) are considered in more detail—or conversely, in which the entire contextualized framework is used to describe the changing dynamics in a historical case or a contemporary situation, possibly from different points of view.

We should point out here that though we draw the framework as four simple areas in two-dimensional space (because it is quick to draw and grasp), this is only a quick reference for discourse, and that simple drawing expands into something more multidimensional in use. For example, we might ask people to consider different aspects of or perspectives on a single situation that might be located in different areas or moving in different directions. In this sense, all sense-making exercises transcend the simple representation of the framework at some point. We consider the framework to be a “jewel of contemplation” that has many facets, like a geographical information system, in which many layers of information are overlaid so that they can be integrated or separated at will; but the ultimate reality is that every representation (including the entire framework itself) is created for a purpose.

The value of the completed contextualized framework lies in two main benefits. First, nearly every contextualization exercise we have seen has ended with expressions of surprise from those participating. They often see, for the first time, patterns that overturn their entrained beliefs about the issue they are considering and about their purpose, goals, and identity. For example, one group completed their Cynefin framework and reviewed it. They had done alternative histories to derive their sense-making items, and they had been asked to provide at least one accident in each fictional time line. On their Cynefin framework, the only items in the chaos domain were those (color-coded) accidents. One participant looked at them and said, “We are being complacent, aren’t we?” The participant meant that they had only been willing to consider wholly chaotic situations when they had been forced to add accidents to the list.

This increased awareness (reflected in many other such stories) is the highest achievement of the completion of the contextualization exercise: that the group should accomplish Descriptive Self-Awareness, or a greater understanding of their own biases and potentials. This is also our goal in helping people go through the process, because it is our place to enable clients to achieve self-awareness rather than to provide “expert” advice, which has a much lower value in practice.

The second benefit of the contextualized Cynefin framework is that it provides a new shared language with which the members of the decision-making group can discuss situations, perspectives and possible actions. This new language is unique to the concerns of the group and abstract enough to cover many particulars, but resonant with meaning so that it tends to be brought up spontaneously when issues are discussed. It can be used to talk about interpretations of current conditions based on gathered data, to evaluate strategic interventions, and to constructively manage conflict and bring about consensus, without removing conflict. Also, multiple groups who have created their own contextualized frameworks can use them as artifacts for negotiation of common meaning. For example, if one group placed “the influence of the Internet on globalization” in complex space and another group placed it in knowable space, both groups would learn something about their respective perspectives. We have been helping groups to use such frameworks to guide discussion among different government branches, for example.

We previously contrasted Cynefin as a sense-making framework with categorization frameworks. We do sometimes use the contextualized Cynefin framework for categorization within a particular context, meaning that the terms used are not generic but specific. Categorization in context has some excellent uses: for example, for training in standard operating procedures, for aligning perspectives and objectives among groups (for example, strategic and operational), or for initiation of new people into the group.

Wider implications

As we mentioned previously, the contextualization exercise is just one of several elements involved in using the Cynefin framework to support people making decisions and crafting strategies. As part of the work we are doing for public-sector and private-sector bodies, including the government, we are design-

ing computer software that uses the Cynefin framework as a sense-making mechanism in tandem with massive narrative databases of world history and contemporary events, in order to support policy-making and decision-making at both the operational and strategic levels in government and industry.

Other elements of our work over the past several years, those having more specifically to do with the use of narrative to enable multi-perspective understanding in the complex domain, are also heavily involved in this work, but we do not have the space

Boundaries are possibly the most important elements, in sense-making, because they represent differences among or transitions between the patterns we create in the world that we perceive.

to describe them here. They are described in recent papers,^{40–41,43} and we plan to feature them strongly in an upcoming book on the subject. Instead, we will delve a little more deeply into some more advanced topics related to using the Cynefin framework.

Cynefin boundaries

Boundaries are possibly the most important elements in sense-making, because they represent differences among or transitions between the patterns we create in the world that we perceive. Depending on how you look at it, any gradient can be a boundary and any boundary can be a gradient, so there is always room for interpretation. We distinguish between types of boundaries on an experiential basis, in the sense of: How does the essence of this boundary, as I perceive it, affect my sense of the situation and of what I should do? In that sense, the boundaries we consider are more like phase changes than physical boundaries (though they could be physical boundaries, if those boundaries coincide with phase changes).

We use a range of metaphor sets to create an understanding of boundaries. One of the most effective is a geographical set, as follows:

1. The *shallow river* can be crossed by anyone at any place, and thus control over crossing is difficult to achieve. However, it is easy to tell when one has crossed it (or when others have) because one's feet get wet. Most of the important events of our lives (e.g. birth, marriage, parenthood, death) are shallow-river boundaries, because anyone can go through them and they "mark" us with the change. Ritual tends to be strongly associated with this type of boundary. An organizational example might be the transition from a new employee to someone who knows the inside stories of the organization. Maintaining shallow-river boundaries is most useful when you want to encourage as much diverse exchange over the boundary as possible, but still retain the capacity to monitor and intervene.

2. The *deep chasm* can be crossed only at bridges, which can be built, demolished, and controlled at will. It is not easy to tell when one has crossed the boundary, but such a marker is not required because only some are allowed through. An example of a deep-chasm boundary might be the front page of a community Web site that requires visitors to register with their names and addresses in order to participate, but does not record individual movements during regular visits to the site. This type of boundary is most useful when the exchange is too important or dangerous to be left open to all possible movements. The selectively permeable barrier found in our cells is a boundary of this type.

3. The *high plateau* is the boundary with the most potential danger, because you may not be aware that you have crossed the boundary until it is too late and you drop off the other side. One of us once volunteered on a mountain rescue team. The worst place to get lost was on a plateau—there are often heavy mists on high plateaus, and people lose their sense of direction and head directly off a cliff. When you are on such a plateau, it is nearly impossible to know where you will come out. (Those fans of certain science fiction programs will recognize this type of boundary in transportation devices, which offer the enticing possibility of putting you somewhere you had not meant to go.) An organizational example of a high-plateau boundary might be the transition that happens during a corporate restructuring, when groups are wandering around looking for support and no one knows who will survive. There are, of course, times when a high-plateau boundary is just what you need: to confuse an adversary, to promote innovation (in limited circumstances), or to disrupt old patterns that have become limiting.

When we are using the Cynefin framework as contextualized by a group to their context and purpose, we ask them to consider the boundaries between Cynefin spaces by using metaphorical types of boundary. We explain that each boundary might have a different form in one direction or the other, and this is where we must break away from the two-dimensional drawing somewhat. In addition, one boundary might have different forms for different people, whose perceptions or circumstances make their experience of the boundary different. Generally, we find three basic levels of sophistication in the use of Cynefin boundaries for sense-making. First, one considers an awareness of *crossing* the boundary, so that one can respond quickly to new conditions after one has arrived on the other side. Second, one considers an awareness of *approaching* the boundary, so that one can sense when change is incipient and respond before the boundary is crossed (perhaps to cross it purposefully, perhaps to avoid it). Third, one considers *managing* the boundary and the perceptions surrounding it, so that one can, for example, put a deep-chasm boundary in place for one's adversary while maintaining a shallow-river boundary for one's own use.

One thing we have found in using boundaries as part of the Cynefin framework is that different people, with different training and personalities, seem to benefit from different uses of boundaries. People who are used to classifying items into categories benefit from removing boundaries, as takes place in the standard contextualization exercise described earlier. However, people who are used to thinking in a more fluid way—about gradients rather than boundaries—seem to benefit more by constructing boundaries than by removing them.

For example, when the workshop is large enough or when we feel that it is required, we often hold a parallel session of contextualization in which people are asked to distribute their sense-making items along a line, ranging from the most tractable items to the most intractable items. After the line has been completed (and there are many negotiations to place items in relation to each other), we ask people to find places along the line at which they feel that the underlying dynamic has shifted. In other words, we ask them to create boundaries along gradients. We then pull the line into a rainbow curve and place it on the Cynefin framework, with the most tractable items in the known domain and the most intractable items in chaotic space. This produces an alternate contextualization, with the same ultimate effect of

creating the Cynefin framework anew, but with a stronger emphasis on the negotiation of where boundaries are found. We sometimes ask people to negotiate boundaries as though they were representatives of the different domains, coming up with a mutual agreement on what the boundary means and where it is placed.

The boundary issue is one on which we continue to pursue a strong research agenda (although of course the entire framework is a work in progress), especially with regard to its use for strategic decision support.

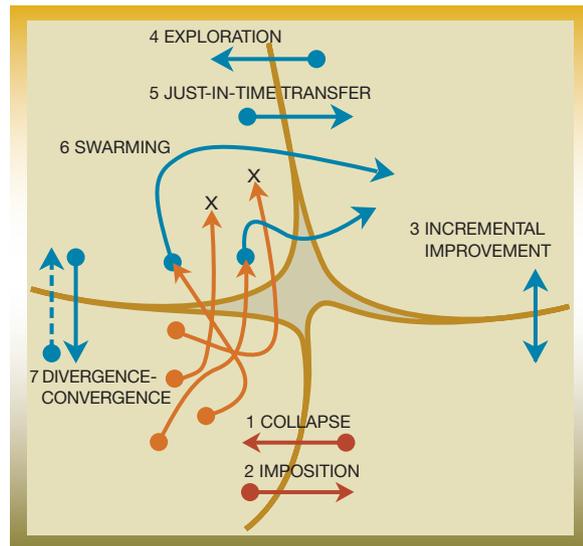
Cynefin dynamics

When people use the Cynefin framework, the way they think about moving between domains is as important as the way they think about the domain they are in, because a move across boundaries requires a shift to a different model of understanding and interpretation as well as a different leadership style. Understanding the differences among the different movements in the framework increases the sophistication of the response of a decision-making group to rapid change. We describe here some of the patterns of movement which we use to help groups consider historical, contemporary, and future change. In general, one of the functions of the Cynefin framework is to increase awareness of the upper domains of the framework and their potential to create sustainable change; several of these movements are designed to make the upper domains more accessible.

Movement at the known-chaos boundary. This boundary is the strongest of the four, in which a perfectly working machine operates inches away from a devastating fire. For that reason, this boundary is the most dangerous—and the most powerful if treated with respect.

Asymmetric collapse (Item 1 in Figure 4) is movement from the known to the chaotic, disastrously. We have seen a tendency for organizations to oscillate between the domains of the known and the chaotic, avoiding the upper domains. Organizations settle into stable symmetric relationships in known space and fail to recognize that the dynamics of the environment have changed until it is too late. The longer the period of stability and the more stable the system, the more likely it is for asymmetric threats or other factors to precipitate a move into chaos. The decision makers in the system don't see things that fall outside the pattern of their expectation, and they

Figure 4 Cynefin dynamics



continue not to see them until finally the system breaks and they find themselves in chaos.

The final stage before the break point is witnessed frequently in history. A good example is the trial of Galileo, in which the Catholic Church accepted that the earth went round the sun for the purpose of mathematic calculation, provided no one said it was actually the case. In retrospect, this was an untenable position, which only delayed and made worse the inevitable collapse. This phenomenon of grasping at order is common in people, governments, academia, and organizations of all shapes and sizes. Often the strongest dominant player in a market will continue with behavior long after its utility, perceived from a different perspective, is exhausted (Boisot⁴⁴ uses IBM as an example of this). Also, senior decision makers and their policy advisors will find ways of fitting reality into their existing models rather than face the fact that those models are outdated, and they will punish dissent (the history of science and business provide examples). Galileo is tried afresh in modern organizations on a regular basis.

Imposition (Item 2 in Figure 4) is movement from the chaotic to the known, forcefully. The consequence of asymmetric collapse is chaos, and the consequence of chaos is frequently Draconian imposition of order, in which the situation is so catastrophic that people accept what would have previously been

unacceptable as the price of order. The problem with this dynamic is that it introduces a new stability that in turn becomes more rigid until the new order breaks in its turn. A familiar example in organizational life is the cyclic reorganization of authority by industry, then by function, then by industry, and so on in an endless cycle; or the fact that well-intentioned revolutionaries sometimes put into place bureaucracies even more stifling than those they overthrew. However, we do not mean to imply that all such transitions are pathological. When order is well aligned with needs, it can bring needed savings and calm. Anyone who has seen a talented teacher take control of a frantic classroom through authority and respect, or a policeman calm a panicked crowd, can understand the utility of imposed yet well-placed order.

Movement at the known-knowable boundary. This is the boundary where the scientific method is believed to operate, though in practice most agree that some un-order is involved in most scientific work (for example, hunches, analogies, networks, local practices, and shared beliefs). This boundary is fluid and permits much traffic as people go about the business of building technologies and pursuing lines of inquiry.

Incremental improvement (Item 3 in Figure 4) is movement from the knowable to the known and back, repeatedly. This type of movement is the best-known and accepted of the types we list here. For many situations, this remains the movement of choice. In a sense, the cyclic flow of information across this boundary is the engine of technological growth. However, it can become pathological if cyclic movements between known and knowable depart ever further from observed reality (as with, for example, the epicyclic models of the solar system or the “science” of phrenology). This sort of movement should be linked occasionally with one of the larger movements we mention later.

Movement at the knowable-complex boundary. The boundary between the knowable and the complex can be a fruitful one for science, and in practice complements the known-knowable border as an engine of new ideas. It is not as permeable as the known-knowable boundary because transitions must translate between two systems of order and from one set of rules to another.

Exploration (Item 4 in Figure 4) is movement from the knowable to the complex, selectively. This movement is often mentioned in the literature on com-

plexity as exploration versus exploitation.¹⁷ Exploration is an opening up of possibilities by reducing or removing central control without a total disruption of connections. In organizations, exploration takes many forms, but trust is key in this movement. One is, in effect, taking a risk by allowing constituent connections to form and strengthen at the expense of central control, and that requires not only good planning and awareness of the “shadow” side

**Informal communities,
which may range from public
to secret in their profile, provide
a rich and fertile source
of knowledge and learning.**

of the organization, but also careful (but unobtrusive) monitoring of the situation. In most organizations there is a strong and often untapped resource to be found in exploratory moves such as this. For example, informal communities, which may range from public to secret in their profile, provide a rich and fertile source of knowledge and learning that is too large and complex to be formally managed. One study of actual practice in IBM Global Services³⁴ indicated some 50–60 official knowledge areas, complemented by many tens of thousands of private areas. By providing spaces in which members of the organization could naturally share with people they trust, a fertile source of learning was created.

Just-in-time (JIT) transfer (Item 5 in Figure 4) is movement from the complex to the knowable, selectively. This movement is often called exploitation in the complexity literature, and it involves the selective choice of stable patterns in complex space for ordered representation. In the IBM example just cited, the 50–60 official communities are able to draw on information coming out of the informal trust-based communities as long as they refrain from disturbing their delicate balance. Knowledge can be moved into the formal space on a just-in-time basis: knowledge is made available when it is needed. Techniques to achieve this include subject-matter flagging and privacy-ensured searching of content. The issues here are those which gave rise to JIT techniques in manufacturing some decades ago: organizations realized that the cost of maintaining stock on the factory floor was out of all proportion to the benefits, with high levels of wastage over and above stock hold-

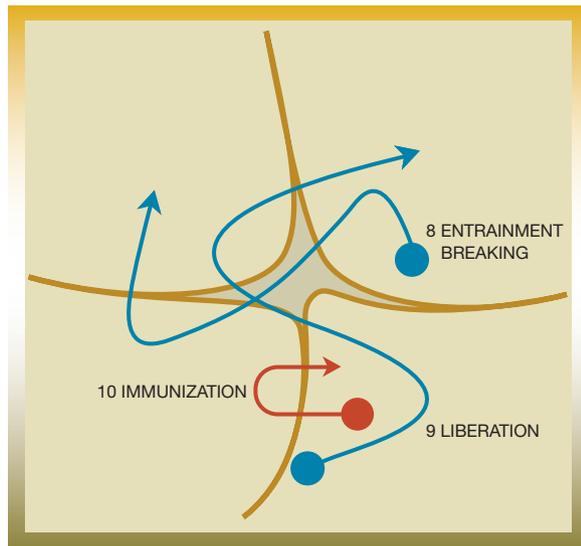
ing costs. In consequence, stock holding shifted back to the suppliers, entering the factory just in time.⁴³

Movement at the complex-chaotic boundary. This boundary, like the known-knowable boundary, is fluid and in fact difficult to delineate. In nature, systems move back and forth across this boundary often. In that sense, traffic across this boundary mirrors that across the known-knowable boundary: one is an engine for technological and scientific order, and the other is an engine for organic order. In the social sphere, we can use the engine of complexity to enable emerging patterns that prove useful.

Swarming (Item 6 in Figure 4) is movement from the chaotic to the complex, to the knowable; first, in an emergent manner and then, selectively. Draconian imposition of order is most appropriate in symmetric conditions and partial remediations, but under asymmetric conditions, or when whole-system interventions are required, we need to move from chaos to the complex, not to the known. The boundary between chaos and order is a chasm difficult to cross, but a vertical transition across the more permeable boundary between chaos and complexity is inherently more manageable. A transition from the chaotic to the complex is a matter of creating multiple attractors, or swarming points, around which un-order can instantiate itself, whereas a transition from the chaotic to the known requires a single strong attractor. For example, if one were trying to evacuate a panicked crowd in a theater on fire, it would make more sense to shout out “the blinking orange lights are above the exit doors,” which is a complex swarming-point trigger that relies on local knowledge only, than to shout out “come towards the back of the theatre,” an ordered trigger that relies on global knowledge which may be unavailable.

After we have achieved the shift from chaos to the complex, then we have the possibilities of many patterns forming around the new attractors; those we find desirable we stabilize through a transfer to the exploitable domain of the knowable; those that are undesirable are destroyed. We have found in several recent engagements that the contrast of *swarming* with *imposition* provides a new language for executives and appears to provide new perspectives on crisis management. We will be studying actual behavior and creating more “subjects” in this field over the next year as our use of dynamic movements within the Cynefin framework develops.

Figure 5 Cynefin dynamics using chaos



Divergence-convergence (Item 7 in Figure 4) is a movement from the complex to the chaotic and back, repeatedly. The active disruption of a complex system to precipitate its move to chaos is less of a change than moving it to either of the ordered domains, and this is easier to manage across a permeable boundary. In knowledge management, for example, informal communities that occupy the complex domain are more resilient when asked to undergo radical disruption in an innovation program than the expert communities of the knowable domain.³⁹ Small start-up companies handle disruption better than large bureaucratic ones, but even within large bureaucratic organizations, there are small groups that can act in the role of start-ups, and they can increase the adaptability of the organization.

Our complexity-inspired workshop techniques, as explained above, make intensive use of the boundary between the complex and the chaotic, in effect cycling between the dynamics of the two states as a sort of pattern generator to create a rich variety of patterns among which to choose—to stabilize and to disrupt—in order to facilitate sense-making.

Visiting chaos

There are some good reasons to move deliberately from order to chaos. There are times when it is necessary to break rigid structures in precipitation of a

natural collapse (as one approaches the boundary), so that the transition can be managed more carefully; and there are times when a strong disruption is the only mechanism that will break up a strong but unhealthy stability. The last three movement types we will consider use the chaotic space for temporary disruption of all connections (possibly within a restricted context) as a stimulant to new growth.

Entrainment breaking (Item 8 in Figure 5) is movement from the knowable to the chaotic to the complex, periodically. In entrainment breaking, we move from the knowable to chaos and thus stimulate the creation of new complex systems as the system rebounds into the complex domain. This is a common approach to disrupt the entrained thinking of experts who, in our experience, tend to be the most conservative when it comes to radical new thinking. The move to complex space is not radical enough to disrupt those patterns; we need to challenge at a more basic level the current assumptions of order. By using the complex space as a staging post, we create a more fertile space of interactions from which we can select stabilization points for the movement to the knowable. A knowledge management example is the creation of formal communities by clustering and swarming informal activities from existing trusted relationships.³⁹ In strategy, this method can be used to create and validate new sources and structures for decision-making.

Liberation (Item 9 of Figure 5) is movement from the known to the complex to the knowable, periodically. Organizations tend to assume that they can design the nature of new systems. For example, an organization that needs new expertise in an area might commission a university to carry out a study, recruit specialist staff, or identify individuals within the organization and assign them new responsibilities. This is a successful and effective strategy when the conditions are suitable for ordered approaches. However, if the situation is uncertain, it is more useful to shift the problem from the domain of the known to the complex. Organizations need to increase both internal and external levels of contact to the point where new patterns can emerge. Boisot⁴⁴ makes the point that companies need to use both hoarding strategies, in which they place defensive barriers around what they know and focus on exploitation, and sharing strategies, where knowledge is shared within and outside the organization with the intent of increasing the volume of opportunities, with the strategic advantage shifting to speed of exploitation of knowledge.

For an organization, sharing strategies involve a process of letting go, of creating freedom within heuristic boundaries to allow new patterns and new leadership to emerge. One of the techniques we use in this area is Social Network Stimulation,³⁹ which aims to stimulate the interactions of agents (or rather identities) within systems to allow the emergence of new coalitions, alliances, and leadership. To use a metaphor, we cast seeds (ideas, deliberately ambiguous goals), which are cheap, across a broad landscape and see where growth occurs. As soon as growth is evident, we respond quickly to shift the newly emergent idea or leaders or coalition into the knowable. We have called this movement *liberation* because it breaks the entrainment of bureaucracy—but like all letting go, it is difficult. This is one of the most threatening of transitions to entrenched managers, but one of the most important.

Immunization (Item 10 of Figure 5) is movement from the known to the chaotic, temporarily. Immunization in chaos is a smaller “visit” to chaotic space that shakes up “the way things are” enough to cause reflection but not enough to destabilize the entire system. Immunization serves two purposes. First, it inures people to the devastating force of chaos so that they will be better prepared to face those forces in the future. A perfect example: it is said that the great director Buster Keaton was able to craft his death-defying stunts (such as a house falling around him, a rescue from a drenching waterfall, amazing pratfalls, and so on) because as a toddler he was lifted out of bed by a tornado and set down unharmed in the street.⁴⁵ Second, immunization brings new perspectives, which cause radical disruptions in stable patterns of thought and lead to new complex patterns. Examples of such events are scattered throughout literature, in the accident that changes a politician’s career, or the chance encounter that causes a lonely woman’s life to fill up with new meaning, or in many other kinds of radical departures that make everything on which one had relied seem meaningless and restricting.

Metaphors are particularly useful agents of immunization because they allow conversation about painful things, enable disruptive and lateral thinking, prevent entrainment of attitudes, and clear out the cobwebs of stagnant ways. One technique we use for this purpose is called the “Grendel game” (material available from the authors, not yet published), which combines anthropological study, complexity theory, and managed war games to create an exciting and innovative learning event. Here, following a study

of the organization, using anthropological techniques developed and proven within earlier research into anecdote collection, a fictional planet is populated by aliens selected to reflect the current culture and new scenarios. This is done with a leading scientist, who in his spare time designs consistent alien environments for use by science fiction and fantasy writers. Members of the organization then seek to colonize the planet in a managed war game. They face their own organization in a metaphorical setting that allows more profound and meaningful learning. Newness is simulated without threat, and the participants are habituated to perspective shift and uncertainty.

Background movements

In any consideration of deliberate change, one must consider what is already going on. The forces of the past tend to cause clockwise drift in the Cynefin space: people living together and sharing mutual needs lead to the emergence of ideas; convenience leads to stabilization and ordering of the ideas; tradition solidifies the ideas into ritual; and sometimes, either lack of maintenance or the buildup of biases leads to breakdown. The forces of the future push dynamics to the counter-clockwise: the death of people and obsolescence of roles cause what is known to be forgotten and require seeking; new generations filled with curiosity begin new explorations that question the validity of established patterns; the energy of youth breaks the rules and brings radical shifts in power and perspective; and sometimes imposition of order is the result. In a sense, these two forces are always pulling society in both directions at once, and this is reflected in organizations as well. The old guard is forgotten at the same time that its beliefs affect newcomers in ways they cannot see. An awareness of these dynamics in the organization must precede any deliberate attempt to affect it by deliberate change.

Use of Cynefin dynamics in practice

Our use of the boundary transitions described here relies on narrative because boundaries are essentially about change and narrative is about change. We may ask people to consider a situation in the past and what movements took place in it from different perspectives, or we may ask people to envision fictional narratives about the past, present or future in which selected movements form the backbone of the story. (Remember this is all done not in the abstract, but using the contextualized Cynefin framework that

makes sense in the context of concern.) These narratives of change are used in the sense-making process, in which they may be quickly created and discarded, or they may be saved and used to generate a shared language about change in much the same way as cycles of folk tales long ago. They are not, however, allowed to stabilize into expectations; they must remain fluid to be useful.

A critical distinction between this type of narrative generation and that of scenario planning⁴⁶⁻⁴⁷ is that the source of Cynefin-based narratives is not a set of expected ranges of expected variables, but a consideration of dynamics in which the variables and ranges are sometimes unknown and perhaps even unspecified. This creates a more diverse, flexible, and changeable set of narratives that should be truly surprising. For example, a scenario-planning exercise done by a police force might result in scenarios exploring the space from high to low crime, corruption, and terrorism. These scenarios are useful in considering a future in which those axes are the most important—but what about the situation where they are overridden by a previously irrelevant factor? Narratives generated using a contextualized Cynefin framework explore spaces in which the dynamic situation arises from any source. Optimally, a wide range of diverse possibilities (chaos from a nuclear accident, a coup, an epidemic, an alien invasion) is best, even though some scenarios might be inconceivable or even nonsensical. Conceivability is not the point: preparation for the unexpected is.

Note that we do not believe these methods should supplant scenario-planning in all contexts, because that method works well in knowable space, where its results are productive. We do believe, however, that these are the methods of choice in complex situations where a wider range of possibility needs to be explored. In fact, as we increase the number of methods available for the un-ordered domain, methods that work well in the ordered domain get even stronger because they are no longer used in situations in which they have limited applicability.

Relationship to other frameworks

We do not pretend that all the basic ideas inherent in the Cynefin framework are new or unique. They can in fact be found floating around history for thousands of years. The distinction between order and un-order (and their interactions) is ancient, as we mentioned, as well as being taken up by recent authors.⁴⁸ The chaotic-complex distinction has been

much debated in recent years, with some saying complexity exists at the “edge of chaos,”¹⁰ some saying that the two phenomena have separate origins and cannot be placed together,^{19,49} and some even saying that the distinction is artificial and arbitrary.⁵⁰ The distinction between known and knowable is widespread and goes back to ancient philosophy. We

As we increase the number of methods available for the un-ordered domain, methods that work well in the ordered domain get even stronger.

do claim originality for the development of the ideas behind the framework in its full form (as we have described it, and in some other aspects outside the scope of this paper) and for the methods we use to make the framework useful in practice (though of course they have their relationships with other action research and sense-making methods).

Courtney’s⁵¹ framework, in which he distinguishes between four states of increasing “residual uncertainty” (uncertainty that cannot be reduced by analysis) questions universal assumptions about the use of known-space methods and tools. He does seem to believe, however, that the level of uncertainty can be “defined”—that is, he has no domain of disorder (which is essentially a state of uncertainty about uncertainty). This, we suspect, may cause people to gravitate to the domain they find most plausible, as we have seen happen. Also, his “toolkits” of choice for each level of uncertainty seem to cover only the known and knowable spaces in the Cynefin framework, with a slight nod to narrative methods in complex space. In effect, he does not break out of the righthand side of the Cynefin framework. He also does not seem to consider the possibility that a single circumstance may contain competing aspects and perspectives with different degrees of uncertainty or that such differences can be used to strategic advantage. Certainly there is much to be gained from all attempts to diversify responses to differing contexts, however.

Returning to assumptions

To complete our circle, we should address the three assumptions we identified at the start of this paper.

As we said at the beginning, these assumptions are valid in some circumstances, but not in others. Let us examine each one again and consider how we can expand these assumptions to a universally applicable set. We will also briefly speculate on emerging applications of the Cynefin framework to management science and practice.

The assumption of order. Relaxing this assumption is the basis of the Cynefin framework, which proposes instead two types of order, each with distinctions inside, and a recognition that uncertainty may exist in distinguishing these types (the domain of disorder). The assumption of order holds for ordered space, the known and knowable, so we need new assumptions for the domains of un-order and disorder. In complex space, we can safely assume that patterns will form, unpredictable in their details but usually recognizable in their basic forms, and that with practice we can learn to detect these forming patterns, stabilize or disrupt them depending on their desirability, and seed desirable patterns by creating attraction points. In chaotic space, we can assume that all connections have been broken, that possibility reigns, that old patterns have been disrupted, and that the outcome is not predictable. In the space of disorder, we know something very valuable—that we do not know. We need to gain more understanding (in every way possible) so that we can find patterns and react to them.

In management, relaxing assumptions of order means recognizing that not all effective solutions are efficient solutions. It does not mean that trust has to be given blindly or that complex processes cannot be affected; it only means that when the means match the context, less energy need be expended for the same result. In a very real sense, managers have successful models available in the domains of un-order in the way they manage their children; they use boundaries and interventions to encourage desirable behavior but do not attempt to control it through goal-based direction. The use of these methods is opening some high-potential developments in managing organizations following mergers and in creating the conditions for continuous learning and innovation.

The assumption of rational choice. Relaxing this assumption means that context and perspective become as important as rationality. This is an important reason that the Cynefin framework is not about “objective” reality but about perception and understanding; it helps us to think about the ways in which

different people might be perceiving the same situation. For example, there is an old folk tale from India⁵² in which a wise man decides that in order to escape an impossible royal demand, he will fake insanity in the king’s court. He is operating in complex space because he is using cultural shorthands to provoke predictable reactions but is gambling that his ruse will seed the pattern he wants to create. He knows that from the perspective of his audience, who are operating in the space where things are bound by tradition and thus known, he appears to be acting chaotically, because they can conceive of no other reason for him to act this way in front of the king (who would surely behead him if he was faking). Thus by proving that he cannot be faking, he pulls off the fake. Understanding not only that there are different perspectives on an event or situation, but that this understanding can be used to one’s advantage, is the strategic benefit of relaxing this assumption. Narrative techniques are particularly suited to increasing one’s exposure to many perspectives on a situation.

In management, there is much to be gained by understanding that entrained patterns determine reactions. This realization has major implications for organizational change and for branding and marketing. Our own work on narrative as a patterning device is gaining presence in this and other areas. Speculating, one of the most significant possible applications of this understanding is a move away from incentive-based targets and formal budgeting processes—both of which, we contend, produce as much negative as positive behavior. It is a truism to say that any explicit system will always be open to “gaming.” Paradox and dialectical reasoning are key tools for managers in the un-ordered domains.

The assumption of intentional capability. The effect of relaxing this assumption is also one of considering context, but more of context on action than on perspective. It means asking not, “What did they have in mind when they did that?” but, “What does it mean that this happened?” John F. Kennedy was presented with such a dilemma when he received two letters in rapid succession from the Soviets during the Cuban missile crisis, one conciliatory and one hard-line.⁵³ Which letter was a wink and which was a blink? It turned out, of course, that both letters were winks but of different actors (Khrushchev and the Politburo, respectively). The different authors of the letters probably saw their actions from the complex space of the Cynefin framework because they understood the internal politics that produced them.

To Kennedy, however, receipt of the two letters plunged him into chaotic space, where he could conceive of no sensible intention to send such contradictory letters—to confuse, perhaps, or to delay while the missiles were being prepared. It was only after scrutiny by linguists that the Kennedy administration was able to understand that the intentions behind the two letters came from different sources. This moved his understanding into complex space where he could respond. It is only possible to consider alternative explanations for actions when one relaxes the assumption that all actions are deliberate.

In management and in strategy, the issue of assumptions about intentions can have pronounced effects. For example, one may assume that the rejection of a new initiative by employees is intentional, when in fact it is an accident of emergent patterns of interaction. Treating the rejection as intentional may create exactly the reaction that the managers wished to avoid in the first place. The same is true of over-reaction to accidental competitor behavior that is presumed to be intentional.

Conclusion

This paper outlines a new approach to strategy, both in policy formation and in operational decision-making. While it is new, this approach also recognizes as critical the value of what has been done to date. As such, it starts to break the fad cycle that has bedevilled management science for several decades. It recognizes the progression of human knowledge, in that something which has provided value is not rendered valueless by new thinking, but is bounded by new insight and legitimized within boundaries, and thereby, made more, rather than less, effective. This approach is, we think, unique in that it recognizes the value and interaction of order and un-order. As such, it also allows us to make a critical distinction between efficiency and effectiveness. Human groups need to be effective; machines and structured human interactions (such as manufacturing processes or the application of rules of engagement) need to be efficient.

Cited references

1. D. Snowden, "The Paradox of Story: Simplicity and Complexity in Strategy," *Scenario and Strategy Planning* **1**, No. 5, 16–20 (November 1999).
2. L. W. King and P. Tice, *Enuma Elish, The Seven Tablets of Creation; The Babylonian and Assyrian Legends Concerning the Creation of the World and of Mankind*, Book Tree Publishers (1999). Available on-line at <http://www.sacredtexts.com>.

3. A. Juarrero, *Dynamics in Action: Intentional Behavior as a Complex System*. MIT Press, Cambridge, MA (1999).
4. I. Asimov, *The Foundation Trilogy: Foundation, Foundation and Empire, and Second Foundation*, Ballantine Books, Inc. (1983).
5. F. W. Taylor, *Principles of Scientific Management*, Harper and Brothers, NY (1919).
6. M. Hammer and J. Champy, *Reengineering the Corporation*, Nicholas Brealey Publishing, London (1993).
7. G. Nicolis and I. Prigogine, *Exploring Complexity*, W. H. Freeman and Co., NY (1989).
8. E. N. Lorenz, *The Essence of Chaos*, University of Washington Press, Seattle, WA (1993).
9. J. H. Holland, *Emergence: From Chaos to Order*, Perseus Publishing, Boulder, CO (1998).
10. S. A. Kauffman, *Investigations*, Oxford University Press, Oxford (2000).
11. C. W. Reynolds, "Flocks, Herds, and Schools: A Distributed Behavioral Model," *Computer Graphics* **21**, No. 4 (SIGGRAPH '87 Conference Proceedings), 25–34 (1987).
12. S. Camazine, J. Deneubourg, H. R. Franks, J. Sneyd, G. Theraula, and E. Bonabeau, *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ (2001).
13. P. Ball, *The Self-Made Tapestry: Pattern Formation in Nature*, Oxford University Press, Oxford (1999).
14. J. Briggs and F. D. Peat, *Turbulent Mirror: An Illustrated Guide to Chaos Theory and the Science of Wholeness*, Harper and Row Publishers, Inc., NY (1989).
15. S. Johnson, *Emergence: The Connected Lives of Ants, Brains, Cities, and Software*, Scribner Book Co. (2001).
16. H. Arrow, J. E. McGrath, and J. L. Berdahl, *Small Groups As Complex Systems: Formation, Coordination, Development, and Adaptation*, Sage Publications, Thousand Oaks, CA (2000).
17. R. Axelrod and M. D. Cohen, *Harnessing Complexity: Organizational Implications of a Scientific Frontier*, Free Press, NY (1999).
18. R. D. Stacey, *Complex Responsive Processes in Organizations: Learning and Knowledge Creation*, Routledge, London (2001).
19. P. Cilliers, *Complexity and Postmodernism*, Routledge, London (1998).
20. D. Zohar, and I. Marshall, *Quantum Self: Human Nature and Consciousness Defined by the New Physics*, Quill/William Morrow, NY (1991).
21. R. Lewin, *Complexity: Life at the Edge of Chaos*, University of Chicago Press, Chicago, IL (1999).
22. G. Weiss, Editor, *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*, MIT Press, Cambridge, MA (1999).
23. M. W. Macy, "Social order in artificial worlds," *Journal of Artificial Societies and Social Simulation* **1**, No. 1 (1998). Available at <http://www.soc.surrey.ac.uk/JASSS/1/1/4.html>
24. R. Axelrod, *The Complexity of Cooperation*, Princeton University Press, Princeton, NJ (1997).
25. R. Dawkins, *The Selfish Gene*, Oxford University Press, Oxford (1989).
26. M. Prietula, K. Carley, and L. Gasser, *Simulating Organizations: Computational Models of Institutions and Groups*, MIT Press, Cambridge, MA (1998).
27. T. Schmalberger, *Dangerous Liaisons: A Theory of Threat Relationships in International Politics*, Ph.D. dissertation, Institut Universitaire de Hautes Études Internationales (1998).

Available at http://www.usc.edu/dept/LAS/ir/cis/cews/html_pages/Dangerous_Liaisons

28. P. Baumgartner and S. Payr, *Speaking Minds: Interviews with Twenty Eminent Cognitive Scientists*, Princeton University Press, Princeton, NJ (1995).
29. D. C. Engelbart, *Augmenting Human Intellect: A Conceptual Framework*, Summary Report AFOSR-3223 under Contract AF 49(638)-1024, SRI Project 3578 for Air Force Office of Scientific Research, Stanford Research Institute, Menlo Park, CA (1962).
30. P. Skagestad, "Thinking With Machines: Intelligence Augmentation, Evolutionary Epistemology, and Semiotic," *Journal of Social and Evolutionary Systems* **16**, No. 1, 157–180 (1993).
31. R. D. Cureton, "e.e. Cummings: A Study of the Poetic Use of Deviant Morphology," *Poetics Today* **1**, No. 2, 213–44 (1979).
32. S. Kostof, *The City Shaped: Urban Patterns and Meanings through History*, Bulfinch Press, Boston, MA (1991).
33. L. Bolman and T. Deal, *Reframing Organizations*, Jossey-Bass Publishers, San Francisco, CA (1997).
34. D. Snowden, "Liberating Knowledge," in *Liberating Knowledge CBI Business Guide*, Caspian Publishing, pp. 9–19 (October 1999).
35. T. A. Stewart, "How to Think with Your Gut," *Business 2.0* (November 2001). Available at <http://www.business2.com/archives/mag/0,1059,317,00.html>.
36. M. Douglas, *Purity and Danger*, Routledge, London (1966).
37. A. Clark, *Being There: Putting Brain, Body and World Together Again*, MIT Press, Cambridge, MA (1997).
38. T. Shea and D. Mavis, Editors, *Managing the Unexpected: Four Case-Study Reports*, U.S. Medicine Institute for Health Studies, Executive Leadership Program Report (2002). Available at <http://www.usminstitute.org/leader.html>.
39. D. Snowden, "Complex Acts of Knowing: Paradox and Descriptive Self Awareness," *Journal of Knowledge Management* **6**, No. 2 (May 2002).
40. D. Snowden, "The New Simplicity: Context, Narrative and Content," *Journal of Knowledge Management* **5**, No. 10, 11–15 (July–August 2002).
41. D. Snowden, "The Art and Science of Story or 'Are you sitting uncomfortably?' Part 2: The Weft and the Warp of Purposeful Story in Business," *Information Review* **17**, No. 4 (2000).
42. H. Beyer and K. Holtzblatt, *Contextual Design: Defining Customer-Centered Systems*, Morgan Kaufmann Publishers, San Francisco, CA (1998).
43. D. Snowden, "Just in Time Knowledge Management," *KM Review* **5**, No. 5, 14–17 (November/December 2002), and **5**, No. 6, 24–27 (January/February 2003).
44. M. Boisot, *Knowledge Assets*, Oxford University Press, Oxford (1998).
45. E. Keaton and J. Vance, *Buster Keaton Remembered*, H. N. Abrams, Inc., NY (2001).
46. P. Schwartz, *The Art of the Long View: Planning for the Future in an Uncertain World*, Doubleday, NY (1991).
47. A. Kleiner, *The Age of Heretics*, Doubleday (1996).
48. M. de Landa, *A Thousand Years of Nonlinear History*, Zone Books, NY (1997).
49. M. Baranger, *Chaos, Complexity, and Entropy: a Physics Talk for Non-Physicists*, Wesleyan University Physics Dept. Colloquium (2001). Available at <http://www.necsi.org/projects/baranger/cee.html>.
50. J. Horgan, *The End of Science: Facing the Limits of Knowl-*

edge in the Twilight of the Scientific Age, Broadway Books, NY (1997).

51. H. Courtney, *20/20 Hindsight*, Harvard Business School Press, Cambridge, MA (2001).
52. A. K. Ramanujan, *Folktales from India: A Selection of Oral Tales from Twenty-two Languages*, Pantheon Books, NY (1991). (The referenced folk tale is "The Boy Who Sold Wisdom," pp. 240–243.)
53. G. Allison and P. Zelikow, *Essence of Decision: Explaining the Cuban Missile Crisis*, (2nd ed.), Pearson Education, Upper Saddle River, NJ (1999).

Accepted for publication April 24, 2003

Cynthia F. Kurtz IBM Global Services Lotus Development, 1 Rogers St., Cambridge, MA 02142 (ckurtz@us.ibm.com). Ms. Kurtz is Principal Researcher for IBM's Cynefin Centre for organizational complexity. She received her B.S. degree in biology from Clarion University of Pennsylvania in 1986 and her M.A. degree in ecology and evolution from the State University of New York at Stony Brook in 1990 (after completing five years of Ph.D.-level graduate work studying the evolution and complexity of social behavior). From 1992 to 1998 she co-owned a small software business, and is the co-creator of two internationally distributed and widely referenced environmental education simulators on plant growth, soil science and evolutionary design. In 1999, she joined the Knowledge Socialization Group at the Thomas J. Watson Research Center, where she carried out pioneering work on narrative in organizations, including the creation of an XML schema for describing stories and storytelling contexts. She moved to the Institute for Knowledge Management in 2001 to work on both narrative programs and complexity programs before helping found the Cynefin Centre in 2002. She is currently leading work on the use of narrative and complexity-informed techniques for gaining multicultural perspective in decision environments, and is mainly responsible for the development of narrative database technology.

David J. Snowden IBM Global Services, 79 Staines Road West, Sunbury-on-Thames, TW16 7AN United Kingdom (snowded@uk.ibm.com). Mr. Snowden is the director of IBM's Cynefin Centre shortly to be based in Cardiff University, Wales. He received his B.A. degree in philosophy from the University of Lancaster in 1975 and an M.B.A. degree from Middlesex University in 1986. His early career was in the specification of decision support systems for international companies, progressing to general management and strategic marketing positions in Dataspace prior to its acquisition by IBM in 1996. He was a director in IBM's Institute for Knowledge Management before founding the Cynefin Centre in 2002. He is an adjunct professor of knowledge management at the University of Canberra, an honorary fellow in knowledge management at the University of Warwick, and has received several awards for his work in knowledge management, including one from the Academy of Management in 2001 for his original contributions to the field.