Puzzles in Organizational Learning: An Exercise in Disciplined Imagination

Karl E. Weick

University of Michigan Business School, 701 Tappan Street, Ann Arbor, MI 48109-1234, USA
email: karlw@umich.edu

Theory construction is in part an exercise in 'disciplined imagination'. One way to discipline our imaginations around the topic of organizational learning, it is to pay special attention to things we forget, values we slight, goals we neglect, facts we avoid and questions we fear. Each aspect of this framework is explored in the paper through a variety of stories and examples. These include Charles Handy's first job, fighting fires, medical diagnosis, Winston Churchill's reflections on the fall of Singapore and the identification of dangerous viruses. These rich and evocative stories offer moments for reflection on underlying issues involved in sense-making, and provide the reader with insights into the relationship between leadership and learning.

Introduction

It is sometimes possible to explore basic questions in the university that are tough to raise in other settings. John Gardner (1968, p. 90) put it well when he said that the university stands for:

- things that are forgotten in the heat of battle;
- values that get pushed aside in the rough-and-tumble of everyday living;
- the goals we ought to be thinking about and never do;
- the facts we don't like to face; and
- the questions we lack the courage to ask.

If we want to develop our thinking in relation to organizational learning then we ought to use the university setting in the ways Gardner suggests. I like Gardner's list because, if theory construction is partly an exercise in disciplined imagination, then Gardner's five qualities are five ways to discipline one's imagination. Essentially, Gardner suggests that if we direct our imagination toward the topic of organizational learning, then we should pay special attention to things we forget, values we slight, goals we neglect, facts we avoid and questions we fear. In this paper I sample learning issues that fall into each of these five categories.

Things that are forgotten in the heat of battle

In their attempts to comprehend organizational learning, scholars sometimes forget to discipline their imaginations consistently with constraints imposed by perceptual and social processes. Both processes, however, can make a substantial difference in our understanding.

Consider what happened to Charles Handy in his first job in an office of 150 people in Kuala Lumpur.

'Ian was the general manager, aged thirty-five – unconventional and earmarked, so they said, for the very top. He had asked to see me first thing. "For your first month here," he said, after greeting me very warmly, "I want you to sit in this room, in that corner there, and to be as inconspicuous as possible. Don't speak when anyone else is in the room, and never, on any account, leave the room, no matter what is going on. You will learn more about this business from watching me for a month than by sitting by some other Nelly, and I will learn too from having to explain to you what is going on and why I did what I did. Feel free, when we are alone, to ask me any questions that are on your mind. And one thing – keep a diary with your impressions, lessons learned, resolutions

made. That will serve as your report on the month – except that I don't want to see it, it will be your report to yourself." It was a fascinating month. I watched a union negotiation, the sacking of a sales manager for dishonesty, the planning of a new refinery – and many other smaller things. I also learned, could not fail to learn, some important lessons in management, the biggest of which was the thrill of being trusted. Some day I said to myself in my report, I will trust some young person as Ian has trusted me.' (Handy, 2000, p. 166)

If we forget that learning is as much perceptual as it is computational, then we look for learning in the wrong activities, and overlook learning in obvious places. To think better is to see better, not to calculate better. Learning and thinking involve seeing through pattern matching, seeing situations as examples of prototypes, observation of the outcomes of mental simulation, developing and maintaining situation awareness, seeing things that others miss, assessing situations for their degree of familiarity, sensing the limitations of frameworks and redoing the frameworks. This perspective on perception and learning comes from Gary Klein's research on recognition-primed decision-making. The centrepiece of this work is the idea that the development of expertise involves perceptual learning, not the accumulation of facts and rules (Klein, 1998, p. 287).

We also sometimes forget that knowledge is not something people possess in their heads but rather something people do together. Some of the best evidence of this oversight is in 'botched handoffs'. Carthey *et al.* (2000), for example, describe adverse medical events that occur in conjunction with neonatal cardiac surgery as the patient is handed over from pre-operative to intra-operative teams, and from intra-operative teams to post-operative care. When the operating team hands the patient over to people in ICU, incorrect information may be communicated (e.g. team says blood loss during the operation was insignificant when it was actually massive, and ICU team is not prepared for related problems); or, no information may be communicated at all (e.g. surgeon starts next case immediately, which leaves ICU staff in the position where they have to take people away from patient care at a key moment in order to get the information they need); or, the surgical team may communicate information too early while the ICU team is trying to set the patient up, and the ICU people fail to hear most of the key information. Recipients never learn what they need to know, and problems compound.

The same failures to learn during handovers happen in wildland firefighting. In 1990, during the 'Dude' fire in Payson, Arizona, six firefighters burned to death when a change of command was botched at 1:00 pm on a hot, windy day with temperatures in the high 90s while the fire was making spectacular runs. The advice that came out of this incident was, 'never hand over a fire in the heat of the day'. What this maxim means is that when a departing crew hands over a fire to an incoming crew, they should do so when it is easiest for the incoming crew to understand what is happening, and to step in and take responsibility. The easiest transitions in fire command normally take place at night when the combination of low winds, high humidity and cool temperatures stabilize fire behaviour and render it most predictable. In the heat of the day, by contrast, the fire is at its most dynamic and most volatile, which makes it harder for the incoming crew to catch up with its rapidly changing character. When there is an attempted handover in the heat of the day, the incoming crew is always behind. It never learns what it faces, its idea of what is happening tends to lag behind what is actually happening and when it fails to get on top of a rapidly developing situation, the level of danger increases dramatically.

The advice to avoid handovers in the heat of the day reflects a shrewd awareness that a relatively small event in the otherwise complex mobilization of resources to suppress a large fire has an important bearing on the success or failure of the mission. It is *not* at all obvious that when a tired crew leaves a fire and is replaced by a fresh crew, things will get worse. The addition of fresh resources normally would be expected to produce redoubled effort and faster suppression. However, this piece of advice alerts people that precisely the opposite could occur.

Values that get pushed aside in the rough-and-tumble of everyday living

Values are an important source of discipline when we frame problems of organizational learning in the imagination. Three values in particular seem

crucial although they are often pushed aside: knowing, respect, happiness.

People value knowing and they learn through public sense-making. The problem is, public sense-making often tends to be done casually in such a way that learning is incomplete. Insecure leaders often flaunt their right to give orders and tend not to say much about why they chose the orders they did. As Paul Gleason said while teaching a class of crew chiefs, 'You can't be threatened by people wanting to know. You're a leader not a wizard?' Public sense-making is an important site for learning research because different forms of briefing and transitions and handovers leave people with very different ideas of what they face and what they need to do. This may sound like nothing more than being sure that people get 'the big picture'. I think that is the wrong image. It is too static. What people are really looking for is 'the big story'. Let me illustrate.

Here is one way to conduct public sense-making. It is a five-part briefing protocol that is being used increasingly often by crew chiefs in the Forest Service to give direction to a firefighting crew.

- here's what I think we face;
- here's what I think we should do;
- here's why;
- here's what we should keep our eye on;
- now talk to me (i.e. tell me if you (a) do not understand, (b) cannot do it, (c) see something I do not).

Here is an example of how that protocol is used. One of the best wildland firefighters in the world is Paul Gleason. His stature comes in part from his work in over 500 serious fires as Crew Chief leading a 19-person Interagency Hotshot Crew (the ZigZag crew). Gleason said that when fighting fires, he prefers to view his leadership efforts as sense-making rather than decision-making. In his words, 'If I make a decision it is a possession, I take pride in it, I tend to defend it and not listen to those who question it. If I make sense, then this is more dynamic and I listen and I can change it. A decision is something you polish. Sensemaking is a direction for the next period'.

When Gleason perceives his work as decision-making, he feels that he postpones action so he can get the decision 'right' and that after he makes the decision, he finds himself defending it rather than revising it to suit changing circumstances.

Polishing and defending eat up valuable time, preclude learning and encourage blind spots. If, instead, Gleason treats an unfolding fire as a problem in sense-making, then he gives his crew a direction for some indefinite period, a direction which by definition is dynamic, open to revision at any time, self-correcting, responsive and with more of its rationale being transparent.

Gleason's example nudges us to think more carefully about what it means to lead when one is thrown into an unknowable, unpredictable context in which the most one can hope for is a plausible direction and plausible updating. Just-in-time learning seems more likely when people stay in motion, have a direction, look closely, update often and converse candidly. This logic derives from the basic process that is involved. That process is embodied in the rhetorical question, how can we know what we think until we see what we say? People need to act in order to discover what they face, they need to talk in order to discover what they think and they need to feel in order to discover what it all means. The 'saying' involves action and animation, the 'seeing' involves directed observation, the 'thinking' involves the updating of previous thinking and the 'we' that makes all of this happen takes the form of candid dialogue that mixes together trust, trustworthiness and self-respect.

What is subtle about all of this is that it is surprisingly indifferent to content. In a way, any old prescription, any old change programme, any old mantra or guru or text will do, as long as that program (1) *animates people* and gets them moving and generating experiments that uncover opportunities; (2) *provides a direction*; (3) *encourages updating* through improved situational awareness and closer attention to what is actually happening; and, (4) *facilitates respectful interaction* in which trust, trustworthiness and self-respect (Campbell, 1990) develop equally and allow people to build a stable rendition of what they face. Whether people become animated because of 'new economic rules', or total quality, or learning organization, or transformation, or teachable points of view, or action learning, or culture change or whatever, they are more or less likely to survive depending on whether their programme engages or blocks these components of sense-making. It is the thrust of this argument, that there is nothing special about the content of change programmes *per se*, that explains their success or

failure. What matters is the extent to which the programme triggers sustained animation, direction, attention and respectful interaction. It is these four activities that make it easier or harder for people to make sense collectively of what they currently face, and to deal with it.

People also value respect. But respect is in short supply for novices. Instead, novices are exposed to testing, hazing, abuse, slamming – all of which make it harder for them to develop skills (see Chetkovich, 2001 for examples of this in the Oakland California Fire Department). In this atmosphere novices are reluctant to speak up, their questions go unanswered, their lack of confidence makes mistakes more likely and trainees who make mistakes are seldom given responsible tasks, which means their learning proceeds even more slowly. At the Mary Pang fire in Seattle, 5 January 1995, at 7pm, the Incident Commander (IC) sent a hose team into the bakery through the front door to keep external fire from entering the building. Because the building was built into a hillside, the IC thought he was sending firefighters in on the bottom floor of a one-story building. In fact, he was sending them in on the second floor of a two-story building. The first floor beneath them was completely engulfed with fire. The fire burned out supports in a wall that was holding up the second floor, the second floor collapsed, and four firefighters fell into the first floor inferno and lost their lives.

One of the more heart-breaking moments of the Mary Pang fire was an incident with a rookie firefighter who was part of the hose crew that was working above the fire on the second floor. The rookie's fire helmet fell off and when he leaned down to pick it up, he felt how hot the concrete floor was. He said to himself, but to no one else, 'I have to remember to ask my Captain when we get back to the station, why concrete gets hot. I didn't know that it did.' We now know that fire was burning under them. And in one sense the rookie knew this. But in another, he did not know what to make of the strange clue, and he did not ask immediately what the strange clue meant. This points to one of the most crucial aspects of sense-making, namely, we need other people in order to do it successfully.

We hear a lot about 'legitimate peripheral participation' as the mechanism by which apprentices learn. The problem is, researchers seldom focus on the dimension of 'legitimacy'. Most of the time they focus on 'peripheral'. Yet a surprisingly large number of occupations treat novices as people who must be tested and who must prove themselves. Legitimacy is earned. And in the earning of it, learning suffers, especially for women and minorities. The idea of 'developing' these novices, supporting their efforts and stepping in before they fail, is foreign to a surprising number of learning settings. What is foreign is the realization that relationships matter in determining the speed and depth of learning.

Finally, people value happiness. But happiness and learning have non-obvious connections. Consider Sylvia Townsend Warner's (1967) famous biography of T. H. White, the author of 'The Once and Future King', which was the basis for 'Camelot'. The biography is built around a vivid pattern in White's life. He was prone to bouts of sadness, melancholia, and depression, which he weathered by following a simple maxim: 'The best thing for being sad is to learn something'. Thus, during his lifetime he raised falcons, headed a school for boys, won a prize for flying airplanes, rode jumping horses in competition, was a deep-sea diver in the era of brass-hat diving suits, learned to sail, was a top-rated hunter and fisherman, drove fast cars, was a painter, did carpentry and learnt how to take shorthand in medieval Latin.

If we put White's maxim in the context of contemporary social psychological research on the effects of happy and sad moods on cognition (Schwarz and Clore, 1996), it suggests some intriguing hypotheses. Studies consistently show that being in a depressed mood facilitates performance on some tasks, especially those that require the detection of co-variation. People in non-depressed moods tend to overestimate the degree of contingency between their own actions and outcomes, and feel that that they are more in control than in fact is the case. Depressed affective states tend to elicit a more systematic and data-driven processing strategy. If a current environment is problematic, then paying attention to the co-variation of acts and outcomes would be adaptive. Thus, when we consider the value of happiness, the question becomes, do people in organizations learn more accurately, when they face a new task in a sad, depressed mood rather than in a positive, optimistic mood? A group that is on the verge of bankruptcy may seem to be good learners under the assumption that they are more motivated to avert disaster and therefore they

will work harder at learning. Maybe, however, motivation has nothing to do with it. Instead, maybe the prospect of bankruptcy induces sadness. And sad people may be better able to learn key contingencies faster and better able to avert bankruptcy by doing the right thing, than is true for people who face bankruptcy with more optimism.

The goals we ought to be thinking about and never do

While there is no shortage of goal statements in discussions of organizational learning, most of the goals discussed focus on efficiency, cost containment, adaptation and speed. The problem is, an imagination that is disciplined by goals such as these tends to ignore the complication that efficient performance is often low reliability performance. While it is true that errors may provide the occasion for learning, and true that low reliability performance is laced with error, it is also true that low reliability settings provide an unstable context for organizational learning due to multiple, simultaneous, ongoing, cumulative anomalies. Imagination disciplined by a richer understanding of the goal of error-free performance, is a precondition for a better grasp of learning that is driven by more commonplace error-laden goals such as efficiency.

Recently there has been an upswing of interest in adverse medical events, medical errors, normal accidents, latent failures and high reliability organizations (Weick and Sutcliffe, 2001). The ultimate goal is to help people learn how to avoid errors or to catch developing errors at an earlier point in time. One fascinating part of this exploration is the renewed attention to the question, just what constitutes an error? Students of learning have a vital interest in answers to this question, since errors and failures are seen as important occasions for learning. If we shift the meaning of an error, then we also shift the meaning of what it means to learn from that error.

Jens Rasmussen (2000) has recently raised some fascinating issues in the context of discussions about errors. Whenever we conclude that someone has made an error, and we search backward for the root cause of the error, we discover that there is no objective stop rule to terminate the causal back-tracking. Instead, we as analysts tend to stop whenever we find an event that we know how to cure. That means that stop rules are arbitrary and so are choices of remedies. Rasmussen proposes a very different idea of what an error involves. He suggests that errors reflect efforts to learn to interact effectively with an environment. When we label something as an error, this often is an 'indication of actors exploring the boundaries of acceptable performance in an unkind environment' (2000: p. 32). Ernst Mach, as early as 1905, saw this possibility, when he said, 'Knowledge and error flow from the same mental sources, only success can tell one from the other' (p. 32). If you look closely at errors in problem solving, they turn out to be the results of attempted solutions. Consistent with that point, experts searching for solutions make more errors than do novices cautiously seeking the tricks of the trade. This phenomenon is readily apparent in Lesgold's (1984) study of training for radiologists and the differences between novices and experts when they size up a complex X-ray. Experts explore the boundaries of an acceptable diagnosis more fully, and discard more trials and construct new possibilities more swiftly than do novices (see Lesgold, 1984, p. 82, Table 1).

Notice that now, rather than study errors, we should focus on strategies to recover from unsuccessful explorations (Rasmussen, 2000, p. 33). When something goes wrong, process operators and surgeons develop and test hypotheses conceptually before they act, to get a sense of consequences. But when do you terminate the conceptual analysis and begin to act? There is no objective stop rule. This means that, from the standpoint of the problem solver, any definition of error tends to be arbitrary. As Rasmussen puts it: 'actions which are quite rational and important during the search for information and the test of a hypothesis may appear to be unacceptable mistakes in hindsight, without access to the situation' (p. 33). So the question of error may well boil down to a very different question, when do I stop thinking and start acting? What makes that question so complex is that we often need to act in order to think.

This is what is so fascinating about the recent interest in improvisation. Some of the interest lies in skill-related issues such as thinking on your feet, making do with fragments and managing the unexpected. But what is also interesting about improvisation is that it poses fundamental issues of action and thought for students of learning. These are the issues that Rasmussen is wrestling with.

Here is a complementary way to frame those issues. When you engage in a standard organizational routine, it is easy to spot errors early in a line of action because the action can be compared with a clear plan or goal. And the error can be corrected. But with improvised actions, you are actually constructing the routine at the very same moment that you are acting it out. You cannot do any comparing of your action with your plan until it is too late to do anything different. An improvisation that is not working usually reveals itself only in its aftermath. When you are improvising, 'an act becomes mistaken only after it has already gone wrong. (But) As it is unfolding, it is not becoming mistaken at all; it is [simply] becoming' (Paget, 1988, p. 56).

For example, think about Rasmussen's example of clinical medicine. Sometimes medical diagnosis is possible only *after* treatment has had some effect on symptoms. Essentially the physician says to herself, because these symptoms responded this way to this treatment, I must have been treating a case of X. Diagnosis is the last thing that becomes clear, not the first thing. Stories about medical errors that come up in conversations between doctors frequently contain phrases such as 'It turns out . . . ultimately . . . subsequent events showed . . . I didn't think it was a mistake then' (Paget, 1988, p. 45). In these phrases, 'back then' means something more than simply back at an earlier point in time. 'Back then' means back when a sequence of treatments was beginning to unfold, and when the physician was acting with uncertain knowledge and with an unknown future coming into being. The line of treatment that the physician was pursuing was not a mistake then. The realization that a line of action is not working occurs only after it has already failed to work. This suggests that occasionally there is a darker side to improvisation. Improvisation that fails reminds us that our understanding always lags a little bit behind our action. When lines of action fail, we need to update, remain resilient and remember that knowledge is not something people possess in their heads, but rather something people do together.

The facts we do not like to face

In a postmodern world, the new 'F' word to be avoided (or flaunted) is 'Facts'. That having been duly noted, it is still true that each of us anchor our explanations in features we treat as 'givens'. Three possible facts that scholars of learning do not like to face are first, that learning involves disbelief; second, that all lessons learned are wrong; and third, that ambivalence is the optimal compromise.

Learning involves disbelieving. The 'old view' of this is that when a statement is presented, we first comprehend it and then later decide whether to believe it or disbelieve it. A 'new view' is supported by Dan Gilbert's (1993) research that suggests comprehension and belief occur simultaneously. Comprehension entails immediate belief. Only later is there the possibility to undo the initial belief and wind up with unbelief. Thus, belief in any new input is the default state; to see and hear something, is to believe it immediately. The implication of this is that if you experience an early interruption while processing a false statement, you will accept it as true because there has been no time to disbelieve it. Fiction should be persuasive since there is immediate belief coupled with no effort to disbelieve. The 'Big Lie' should work if you can keep people from counter-arguing. When people engage in what they call 'reflection', the bulk of this activity may involve unbelieving things they initially believed to be true. Reflection that consists of long stretches of beliefs that unravel, may be part of what contributes to melancholia, depression and cynicism such as we saw in the case of T. H. White. If comprehension is synonymous with belief, then the value of experience may be overrated. This follows since both true and false inputs are believed equally. When people then act on this experience, they are misled by the false inputs unless they have spent time sorting through their experience.

A second disturbing fact is that lessons learned are always wrong. I think hindsight causes us a lot more trouble than we realize. Hindsight and second-guessing is a huge problem for incident commanders at urban fires. The fire chief in Phoenix put the issue this way: 'The number of faults in a fire operation is in direct proportion to the number of viewers; the intelligence of the viewers is in direct proportion to how late they arrive' (Flin, 1996, p. 22). The chief also said, just before the remark I quoted, that, 'considering what was going on at the time a decision was actually made will many times effectively refocus 20-20 hindsight'.

You cannot learn a lesson of any kind until some outcome has occurred. Most outcomes induce selective hindsight. If hindsight produces severe editing of the causal chain that led up to the outcome, then the lessons we feel we learned probably never happened. The sequence is trial, error, hindsight, uncover bad precursors, ignore good precursors, change all or most bad precursors (even those that were okay and were neutralized or undone by other events), try again and fail again, but now for different reasons.

The trick in an after-action review is to break the stranglehold of hindsight. I continue to be impressed with the way Winston Churchill questioned one of the darkest moments in his life. During WWII, Churchill made a colossal error when he failed to realize how vulnerable Singapore was to attack by a Japanese land invasion. This error led to Singapore's downfall. After the collapse Churchill asked four questions: 'why didn't I know', 'why wasn't I told', 'why didn't I ask', 'why didn't I tell what I knew' (Allinson, 1993, pp. 11–12). Those four questions are questions of sense-making. Those four questions take seriously the idea that knowledge is not something people possess in their heads, but rather something people do together. Those questions are also sufficiently different that they could uncover both correct and incorrect antecedents of his decision.

The final troubling fact, made clear by social evolutionists, is that ambivalence has survival value. If ambivalence is the optimal compromise, then what do we need to learn? Almost any interpretation of a dynamic, changing, uncertain problem may be plausible, in the sense that it explains some of the attributes of the problem. But almost any interpretation is also flawed and incomplete in the sense that many other potential features are simply ignored. This state of affairs can be managed if a system acts as if the achieved interpretation is both correct and incorrect. People treat the diagnosis as correct in order to start some activity that will at least address the problem. But they also remain alert to the possibility that their interpretation may be wrong. Acting as if they are right changes the problem, generates information and supplies a framework within which whatever happens has some meaning. Acting as if they are wrong promotes doubt, updating and revision of the initial impressions. This is the sense in which ambivalence is the optimal compromise. Donald Campbell, following William James's lead, spells out the logic: 'In multiple-contingency environments, the joint presence of opposing tendencies has a functional survival value. Where each of two opposing tendencies has survival relevance, the biological solution seems to be an ambivalent alternation of expressions of each rather than the consistent expression of an intermediate state. Ambivalence, rather than averaging, seems the optimal compromise' (Weick, 2001, p. 377).

People forego the advantage of ambivalence when they ignore the flawed quality of an achieved interpretation, seek confirmation of their initial diagnosis and continue to ignore non-fitting cues. The non-fitting cues remain in play even though they continue to be undetected. Eventually they may surface in the form of an advanced problem that is much harder to contain. The decision to ignore doubts may make sense if there is a premium on speed, if a reputation for decisiveness is at stake, and if one has not yet fully understood what it means to make sense of something novel that is 'emerging'.

Ambivalence can be a property of networks and distributed cognition, as well as a property of individuals. This is well illustrated by the case in 1999 where the Centers for Disease Control (CDC) mistakenly diagnosed an outbreak of the West Nile Virus (WNV) as an outbreak of St Louis Encephalitis (SLE). Flushing Hospital in New York City asked for CDC's help on 2 September in analysing serum samples from patients who had a strange set of symptoms. CDC's Vector-borne lab concluded the next day that SLE was 'the most likely cause' for the patients' illness (Minority Staff Report, 2000, p. 9), that they had screened for antibodies to 'common American arboviruses that cause encephalitis' (p. 9) and that their screening 'has ruled out several other viruses' (GAO, 2000, p. 45). It is true that CDC had ruled out viruses other than those in the flavivirus family. But, since flaviviruses cross-react in antibody tests, a positive reading such as they obtained meant that any one of several viruses could have been infecting people.

It took three weeks to discover that SLE was a swift but inaccurate diagnosis. There are several explanations for the delay. First, SLE had occurred in the United States before, but it had never appeared this far east (GAO, 2000, p. 10). If there is an unexpected event, then that surprise

absorbs most of your attention. If, improbably, there is a second unexpected event such as the unprecedented appearance of an Old-World virus in a New-World setting, then that second surprise may be missed altogether. Second, the Vector-borne lab shifted from acting as a reference laboratory to acting as a testing laboratory soon after SLE was announced. The Vector-borne lab was inundated with requests to test serum samples from NYC for the presence of SLE (GAO, 2000, p. 47). In accepting this assignment, the lab shifted from a learning orientation to an execution orientation, and did virus-specific assays in which they looked specifically for what they presumed they would find, namely, SLE. The question was not, what do we face here? The question was, is SLE in this sample? This is the mindset of confirmation, not the mindset of discovery and accuracy.

Discovery occurred elsewhere and was fostered by a context of ambivalence, doubt, questions about the accuracy of the SLE diagnosis and openness to the possibility that an upswing in animal deaths (e.g. horses, crows, exotic birds in Bronx zoo) might be evidence of a single virus that affected humans as well as animals. The learning that prompted a reversal of the initial diagnosis can be understood as the activation, over time, of a network that gradually came to doubt the SLE diagnosis. As the network grew larger and more heterogeneous, it connected with more pockets of doubt – present mainly in the animal-health community – which meant that the network as a whole became more ambivalent. Some portions acted as if there was an outbreak of SLE while other portions simultaneously acted as if there was an outbreak of something that resembled SLE but was not SLE (e.g. New York State sent samples to a lab at UC Irvine because it made a different set of assumptions than did CDC [GAO, 2000, p. 11]). The growing incidence of interorganizational ambivalence enacted a collective mind that deployed more intelligence toward the samples being tested, and was less committed to finding justification for the original announcement. Networks that embody ambivalence may be better forms for swift accuracy in a context of emerging diseases than networks that embody certainty. Stated less structurally, networks capable of organizing for ambivalence learn faster than do networks that are organized solely for belief or doubt.

The questions we lack the courage to ask

The beauty of disciplined imagination is that any question can be fielded, and any answer can be mentally simulated. Nevertheless, the labels we use to frame the question are themselves protected from scrutiny. And therein lies the challenge for a scholar of organizational learning. I conclude with a very brief example of a question some lack the courage to ask.

Can uncertainty be reduced? Dietrich Dorner (1996) has raised the disturbing possibility that when people experience uncertainty and gather information to reduce it, this often backfires and uncertainty increases. As a result, it is those with the least information who are most certain. The more information that is gathered, the more doubts accumulate about any option, and the more paralysed the individual may become. Given all the hype suggesting that information is good and more is better, one can well imagine that many will reject this so that they can preserve a simpler picture in order to remain in action. If those data that are rejected are the very data that disconfirm an observer's expectations – and those are the data that people are most likely to avoid – then this makes it more likely that the system in which they are embedded is incubating a deteriorating situation that could explode.

Whether it is a system or a person that is unable to tolerate the discomfort of information gathering, the result is the same – heightened vulnerability. More importantly, the condition of heightened vulnerability is invisible. There is no felt need for learning, right at the very moment where learning is most urgently needed. The mantra 'keep it simple stupid' (KISS) may in fact mean, keeping it simple is stupid because it induces stupidity. It takes students of learning to spot the possibility that 'kiss' may have the appearance of great strategy but the reality of terrible practice, and then to champion doubt, complication and wary simplification.

Conclusion

In a world where there is no shortage of questions and puzzles, an essay like this one, which adds more questions and puzzles, can be a mixed blessing. But, if you re-examine the puzzles I have imagined and the categories that I have used to discipline that

imagining (memories, values, goals, facts, questions), you will see that privileged concepts and assumptions have been dropped, retained ideas have been held more lightly and newer proposals have been stripped of weighty dogma. All of these moves presume that scholars of learning should carry their hypotheses lightly and be willing to drop heavy tools in order to become more agile theorists. The idea of dropping one's tools to regain lightness and agility is old news, dating back at least as far as Lao Tzu (Muller, 1999) who said,

> 'In pursuit of knowledge, every day something is acquired;
>   In pursuit of wisdom, every day something is dropped.'

At a time when organizational learning is usually treated as an issue of acquisition, it is easy to forget that learning is also sometimes an issue of dropping in the interest of wisdom.

References

Allinson, R. E. (1993). *Global Disasters: Inquiries into Management Ethics*. Prentice-Hall, New York.

Campbell, D. T. (1990). 'Asch's moral epistemology for socially shared knowledge'. In I. Rock (ed.), *The legacy of Solomon Asch: Essays in cognition and social psychology,* pp. 39–52. Erlbaum, Hillsdale, NJ.

Carthey, J., M. de Leval, J. Reason, A. Leggatt and D. Wright (2000). 'Adverse events in cardiac surgery: The role played by human and organisational factors'. In: C. Vincent and B. De Mol (eds), *Safety in Medicine,* pp. 117–138. Pergamon, Amsterdam.

Chetkovich, C. (1997). *Real Heat: Gender and Race in the Urban Fire Service*. Rutgers University Press, New Brunswick, NJ.

Dorner, D. (1996). *The Logic of Failure*. Addison-Wesley, Reading, MA.

Flin, R. H. (1996). *Sitting in the Hot Seat: Leaders and Teams for Critical Incident Management*. John Wiley, New York.

GAO (2000). *West Nile Outbreak: Lessons for Public Health Preparedness.* (Report No. GAO/HEHS-00-180.) United States General Accounting Office, Washington, DC.

Gardner, J. W. (1968). *No Easy Victories*. Harper, New York.

Gilbert, D. (1993). 'The assent of man: Mental representation and the control of belief'. In: D. M. Wegner and J. W. Pennebaker (eds), *Handbook of Mental Control*, pp. 57–86. Prentice-Hall, Englewood Cliffs, NJ.

Handy, C. (2000). *21 Ideas for Managers.* Jossey-Bass, San Francisco, CA.

Klein, G. (1998). *Sources of Power*. MIT Press, Cambridge, MA.

Lesgold, A. M. (1984). 'Human skills in a computerized society: Complex skills and their acquisition', *Behavior Research Methods, Instruments, and Computers*, **16**(2), pp. 79–87.

Minority Staff of Senate Governmental Affairs Committee. (2000). *Expect the unexpected: The West Nile Virus wake up call (Report to Senator Joseph I. Lieberman).* Minority Staff of Senate Governmental Affairs Committee, Washington, DC.

Muller, W. (1999). *Sabbath: Restoring the sacred rhythm of rest*. Bantam, New York.

Paget, M. (1988). *The unity of mistakes: A phenomenological interpretation of medical work*. Temple University, Philadelphia.

Rasmussen, J. (2000). 'The concept of human error: Is it useful for the design of safe systems in healthcare?'. In: C. Vincent and B. De Mol (eds), *Safety in Medicine*, pp. 31–48. Pergamon, Amsterdam.

Schwarz, N. and G. L. Clore (1996). 'Feelings and phenomenal experiences'. In: E. T. Higgins and A. Kruglanski (eds), *Social Psychology: Handbook of Basic Principles*, pp. 433–465. Guilford, New York.

Warner, S. T. (1967). *T. H. White: A biography*. Viking, New York.

Weick, K. E. (2001). 'The attitude of wisdom: Ambivalence as the optimal compromise'. In: K. E. Weick, *Making sense of the organization*, pp. 361–379. Blackwell, Oxford.

Weick, K. E. and K. M. Sutcliffe (2001). *Managing the unexpected*. Jossey-Bass, San Francisco.