

Stories from the frontier

David Snowden

Cynefin Center for Organizational Complexity, UK

It's a sad reflection on the tyrannies of modern life that one of the few opportunities one gets for uninterrupted writing is on a plane. This piece comes to you courtesy of an upgrade on a flight to Washington, the start of five week sequence of flights to either sell, check-up on, or take part in a series of complexity-based engagements. In-seat power, a bad film and a Gin & Tonic now stimulate an entrained pattern of concentration! It may even be necessary to install a second hand airline seat at home to achieve the same results! Of course it will also need a small screen DVD, an endless supply of romantic comedy and past episodes of *Everyone loves Raymond*. Entrained patterns are interesting. It was the same with my thesis, written to a constant cycling of a recording of the Boulez Ring. To this day it means that the right leitmotif can trigger a sustained period of panic induced synthesis of partially understood ideas. Patterns can be stimulated, but they can also blind. The failure to detect weak signals in a market, in employee reaction to change, or in the field of counter terrorism can result in catastrophic failure or missed opportunities. The patterns of past failure and success determine current perception.

I have argued on many occasions that emergent processes generally, and in the context of narrative in particular, provide a means to disrupt entrained patterns of individual and collective cognition. They also enable new perspectives and critically increase the scanning range of humans for weak signal detection. Now this is something I have always believed, have practiced and have achieved anecdotal success. It has also passed the ultimate test: people come back for more. However in the last month things have moved on. I have had the privilege of working with Gary Klein (famous for *Sources of Power* and founder of the naturalistic movement in decision theory) to experiment with different groups of military and homeland defence decision-makers over a two-week period. The aim, in a properly controlled experiment, was to see the degree to which a synthesis of our two approaches could improve weak signal detection and scanning capability. I am pleased to say that the experiments

succeeded and the results will be published at some stage in the not too distant future.

I think it is highly significant that this experiment was commissioned as it means the complexity-based ideas are starting to shift to a more serious plane, requiring experimental confirmation of effectiveness. That assumption is supported by the current round the world trip, which is more about checking up on a series of engagements, rather than selling to sceptical audiences (although those are still there). In my judgement the tide seems to be turning, approaches based on complexity are gaining acceptance and not just among the early adopters. We are also moving from small-scale experimental projects to more extensive adoption. However with increasing adoption comes danger. The last few decades have seen a succession of management fads, some good, and some bad, none of which have fully delivered on their promises. All of them are based on a cause and effect model of the *ideal*. The depressing thing is that a lot of complexity consultancy is following the same model. Just as traditional management theory confuses *correlation* with *causation*, so too much complexity work confuses *simulation* with *prediction*.

Naturalistic vs. ideal approaches

All current management approaches work from the *ideal*. Both the mechanical metaphor of Business Process Re-engineering, and its recent manifestation in the anal retentive approach to measurement (when carried to excess) known as Six-Sigma, start with a definition of how things should be, proceed to a gap analysis with how things are and then seek to close the gap over a defined time frame. Systems Dynamics, with its metaphor of the human body (the DNA of the organization, etc.), also works from the *ideal*, although it is more based on behavior than process in its popular manifestation: defining a set of values for employees and then initiating change processes to attempt to imprint those values as 'mental models'. Like Process Management it's a valid approach, but it has limits. It can only be applied to systems where there are known or knowable repeatable relationships be-

tween cause and effect. Only in these circumstances can we identify the necessary steps to close the gap between *reality* and the *ideal*.

The alternative to approaches based on the *ideal*, is to take a naturalistic approach. Here we start with how things are and then we attempt to discover reasons why, ideally using material from the physical sciences. If we have a scientifically valid reason, then we can scale the naturalistic solution. For example, the first fit pattern match that characterizes human decision making was discovered by Klein and others experimentally. We now know from science, that this comes from multiple patterns stored in the long-term memory and sequenced in frequency of use. Now we can aim to manage those patterns, rather than trying to enforce a 'rational' model based on an information-processing model of the human brain. So the naturalistic approach starts with how things are, then seeks to stimulate or manage the evolution of patterns - disrupting the negative and reinforcing the positive - to allow the discovery of futures that we could not anticipate, but which are better (or at least more realistic) than the ideal. Naturalistic approaches are (sic) natural to complexity thinking and it allows complexity practice to draw on, and inform an established tradition.

The application of complexity should work in the same way. We have an approach and a developing set of methods and tools that are based on concepts such as emergence and coevolution. To apply the structures of outcome based *idealism* and the associated industrial recipes of consultancy practice would be at best a mistake, and at worst a betrayal. It's all too easy to see this happening. There are many examples: the use of complexity language and concepts within a systems dynamics framework for example is common; the attempt to use computer-based simulations to go beyond augmentation to replacement of human decision-making is another. Which is not to deny the value of either, but to argue for the need to treat human complexity as a new and fascinating subject, for which the entrained patterns of management science research and management consultancy practice are no longer appropriate. Properly understood we can transform existing processes by focusing on the starting conditions and managing evolution, rather than attempting to constrain natural processes within the strait jacket of the *ideal*.

Weak signal detection

To return to the work with Gary Klein which was for me the highlight of the month, both for the creation of a proper controlled experiment to validate the use of complexity ideas, and for the chance to work with one of the giants of decision theory. When you get such a project, the client works you hard, so in addition to running the experiments we were frequently shipped out to present to various military audiences interested in our work. The presentations and the conversations were as valuable as the experiments. One of the things that we discovered we had in common is that we had both started from practice, and then moved into academia, and that both of us were convinced that had we stayed in an academic track then we would have been unable to achieve the same degree of invention. It was the interaction of multiple academic concepts with practice that had produced, its coevolution that had produced the ideas, not the study of practice through the lens of hypotheses (this is my own summary of the discussion by the way, Gary has no liability for the opinion expressed therein).

One preliminary conclusion from those experiments builds on the narrative theme of the last edition. It seems that narrative forms of knowledge capture show a higher capacity to pick up and detect weak signals than structured conversations and analysis. This was one of the first occasions on which we have used the general body of narrative capture and self-indexing software to act as a recording device in a structured experiment and it worked well. At the end of each experimental run (we had three for each command group in six experiments in all) we did a basic direct questionnaire to check what the group and individuals had picked up. In addition we placed the command group in different contexts (a briefing to a senior commander, a briefing to a journalist, and a grand father's stories to his grandchildren entering the service) and prompted them to tell the story of what had taken place. The first approach lacked ambiguity, the second (like all narrative techniques) features ambiguity. Both were useful, but the richness of narrative material was a distinguishing feature.

Of course picking up a signal is just part of the problem. The issue of weak signal detection in sense making has at least four aspects:

1. Scanning: Did we pick up the signals at some level?

2. Perception: Did those signals register in some way?
3. Attention: Did we see something as worthy of our notice?
4. Response: Can we act (or get other people to act) on those signals?

These are separate and distinct issues. If we look at disaster, and other, reports they generally show (with the benefit of hindsight) that the necessary information was available somewhere. The argument then follows that we should have *joined up the dots* and any failure is one of knowledge management or organizational structure; hence the common solution to any disaster is a major reorganization and the purchase and installation of a new computer system. The preliminary report on 911 argued that there was a failure of knowledge management. Around the same time a report came out on the shuttle disaster. It argued that all the right information had been passed to the right people, but they had not paid attention to it. It was of course the second shuttle disaster; the first had reached a similar conclusion to that of the preliminary report on the build up to the events of 911. The connection between the two is not often made.

It links with a general error in *idealistic* approaches to decision theory. It is not enough to have the right information in the right place at the right time for the right people. Yes, we need information (or data or signals, but let's not go there in this article), but there should be no assumption that any of the four aspects of sensemaking above will necessarily follow from the availability of the said information. Just think of the sheer volume of data and email that arises in a risk adverse organization. If you send an email, or write a report that argues that there may be a disaster you have nothing to lose. If nothing happens everyone will forget what you wrote (if they even paid attention to it in the first place). If something does, then you will be vindicated. For the decision-maker the sheer volume of information and advice is excessive to the point of dismissal. Most senior decision-makers therefore filter their attention to information provided by trusted sources - those who have been successful in the past - but that is not necessarily a predictor of the future.

Scanning patterns not information

Now there is a famous short film of a group of white and dark costumed basketball players (www.viscog.net) passing a ball. If you instruct the group to count the number of times the ball is passed by those with white shirts, then a range of answers will be given by the group. Rarely if ever does anyone notice a person in a black gorilla suit walk to the centre of the group, beat his/her chest and then walk out of shot. They are paying attention to counting the ball and they miss the gorilla. Of course if they see the gorilla there is little chance they can accurately count the number of passes. All the information is present, but the context of your observation determines what you scan for. I see that as an argument for increasing the diversity of the scanning capability within an organization, and thereby an argument against widespread uniformity. If different people, or groups of people, have different perspectives then the group as a whole has different scanning capabilities.

Now those who see weak signal detection and related problems of sensemaking as information processing, focus on computer-based massive data processing on the basis that more data will lead to an increased capacity to detect weak signals. The other common argument is the need for *common ground* (or in the case of computer systems, common taxonomies) between different agencies or silos to ensure that they understand information in the same way. This comes from the information processing models that have dominating thinking about decision-making. One of my daughter's (currently eighteen months away from going to University) psychology textbooks states, "We work from the assumption that the human brain is a limited capacity information processing device." It is an appalling statement, but one that still dominates much theory and far too much practice in management.

The problem here is that humans are pattern-based intelligences. I remember from when I was involved in forecasting software development many years ago, it didn't matter how accurate the computer was, if humans could not understand the basis of the forecast they would second guess it. Humans evolved for a complex world in which prediction is dangerous, and hedging against uncertainty a sensible strategy. If you understand, albeit with no precision, the basis of something, then your capacity to act quickly to adjust or change direction

is enhanced. If you are following a computed prediction then you are in danger. We need to build our own patterns, in order to act.

The argument for *common ground* means all of us will be counting the balls, and no one will see the gorilla. We have conducted experiments which show that increasing bias in a group, and more importantly ritualizing dissent (the ritual prevents it being personal and this is very different from the concept of a devil's advocate) increases not only the ability of a group to scan a range of data, but also to pay attention to it and agree actions that mitigate risk. To manage this we need to create ways of representing an uncertain environment that will permit action to be taken without knowledge of outcome. I will deal with this very briefly in the final section of the piece, but return to the theme later in the year.

The criticality of context and uncertainty of future outcome

Representation of a complex space requires us to recognize the usefulness of ambiguity in avoiding pattern entrainment. One of the techniques we developed for strategic mapping is a perspective model represented in the mnemonic ABIDE. It represents the things we can manage in a complex space, namely Attractors, Barriers, Identity, Dissent and Environmental factors. We have used this on several occasions and one of the key findings is that the mnemonic works where the future is uncertain, but fails when it is known. Last year, working with some serious players with real experience in anti-terrorism, we attempted to use the mnemonic to retrospectively understand the events leading up to 911. It proved to be near impossible as the discussion was taking place several years after the event, and the group could not avoid the "no we should have known" statements of retrospective coherence. So by accident over coffee, a conversation developed about the then developing events in Falluja, Iraq where the outcome was unknowable. The group picked up on the earlier exercise and had no problem describing the situation in terms of attractors and barriers. The technique worked where there was uncertainty. We found the same in our recent experiment with Klein. In previous work with the client the idea of ABIDE was attractive, but people found it an abstract concept. However in the experiment, in the middle of a real simulation with high uncertainty, it was fairly easy for the alternative

description to be used, and for that description (i.e., what are the starting conditions not the desired end point) to increase the capacity of the group to take multiple factors into consideration.

This crystallized for me what has been an emerging understanding over the last few years, which is to say complexity work is contextual in nature. In the context of uncertainty they are intuitive and easy to understand. In the abstractions of a training course, or a discussion, they are difficult to grasp. Theory coevolves with, but does not arise from, practice. It is also true that in practice, when things are uncertain, seemingly difficult concepts become less so, as the old patterns of thinking (which made them difficult to understand) are no longer sustainable, so new models become acceptable and useful.

The bit on measurement (an apology)

In the last edition I provided examples of a complex systems approach to allowing the emergence of meaning from narrative. I contrasted this emergent process with either managerialism (which some readers may recall I had been accused of) and the tyranny of experts who seek to deconstruct and interpret narrative supposedly to remove ideology, but in practice use the pattern entrained filters of their own expertise to impose an ideology. I promised that in this article I would continue the narrative theme by looking at quantitative approaches derived from narrative practice to look at issues of impact and change. In practice the deadline for this piece has come round two months earlier than expected, so that will wait for the next edition.