

Liberating Knowledge

Understanding the sense making communities in the complex ecologies of the modern organisation

“Some people do not become thinkers simply because their memories are too good”

Nietzsche

Currently there are two distinct schools of thought in knowledge management. These schools operate according to different metaphors of the organisation and the society in which it operates.

The first uses a mechanical metaphor. Here the organisation is seen as something that, with sufficient study and analysis, can be understood and prescriptive models can be created that will produce consistent and beneficial behavior. This school has an honorable tradition, originating in the work of Frederick Taylor and other founding fathers of scientific management in the early part of the twentieth century. It culminated in business process reengineering, itself having roots in cybernetic theory. This approach has also been driven by the significant growth of consultancy over the past 30 years. The financial model of most major consultancies requires extensive intellectual capital reuse. Once created, a method can be rolled out consistently in a variety of organisations by junior staff. One partner can now support far more consultants than was the case even five years ago. The business schools in turn have geared themselves to the production of analytically focused and ambitious graduates to feed this demand. However, the scientific model is that of Newtonian physics and, like Newtonian physics, it is no longer good enough and is certainly not universally applicable – although it is still useful in the majority of day-to-day circumstances.

The second school operates from an organic metaphor, seeing each organisation and its environment as a unique complex ecology comprising multiple inter-dependent and inter-causal units. The organic school recognises the historical value of the mechanical metaphor in creating the modern organisation and driving the efficiency improvements of the last few decades. The distinction between mechanical and organic metaphors is not new. It first occurs in the sixties (Burns and Stalker, 1961) and is foreshadowed in a collection of essays looking at representation

“...Neither is new wine put into old wineskins; if it is, the skins burst, and the wine is spilled, and the skins are destroyed; but new wine is put into fresh wineskins, and so both are preserved.”

and anti-representation, which included some of the early work on knowledge and autopoiesis – best understood from its Greek root *auto* self plus *poietein* to make, produce, remake, conceptualise (von Krogh and Roos, 1996).

benefit has been painfully wrung from the mechanical metaphor. Indeed, in some

The mechanical metaphor will continue to be useful in quality management, process improvement and system design along with other structured and known aspects of an organisation’s operations. However, with BPR, the last drop of

implementations process reengineering often went a step to far and removed the necessary redundancy that enabled organisations to respond organically to change.

In a new age of uncertainty we need more flexibility and responsiveness than can readily be provided by a mechanical understanding of the firm: the ability to sense change on receipt of incomplete or partial data or stimulus; to respond in such a way as to reduce uncertainty for your organisation, but not for your competition; to make your organisation lucky by being in the right place at the right time. These are the survival characteristics of the modern resilient organisation.

The underlying issue here – and the one that the organic school seeks to understand – is the increasing level of uncertainty in the business environment. The boundaries between supplier and customer are blurring. New markets and new market leaders are seemingly created overnight with little investment. Information warfare techniques are practiced between competitors. Key staff and their teams sell their intellectual assets to the highest bidder. We live in an era of change and uncertainty, in which the ecology is not sufficiently stable for any mechanical model to approximate to reality.

The impact of knowledge management

The growing practice of knowledge management has driven the move away from mechanical models. Organisations have recognised the importance of actively managing their intellectual capital. With this recognition came the distinction between tacit and explicit knowledge, popularised by Nonaka and Takeuchi (1995), although it originates in the many works of Polyani earlier in the century.

While explicit knowledge and information had been the subject of successful mechanical approaches in the past, the human, tacit community-based element was not susceptible to such an approach. Many companies focused on document repositories and collaborative sharing on the assumption of rational behaviour that was not achieved in practice. It is not possible to build an intellectual capital management system (ICMS) if we see it as a universal application. Intellectual capital is too diverse, too complex and too heavily dependent on individuals and communities who do not behave rationally. Neither do we want them to behave rationally – to do so would drive out innovation and relationships, both of which are skills for the knowledge economy.

In the CBI Guide to Knowledge Management, I defined knowledge management as *the identification, optimisation and active management of intellectual assets, either in the form of explicit knowledge held in artifacts or as tacit knowledge possessed by individuals or communities. The optimisation of explicit knowledge is achieved by the consolidation and making available of artifacts. The optimisation of tacit knowledge is achieved through the creation of communities to hold, share and grow the tacit knowledge. The active management of intellectual assets is the creation of management processes and infrastructure to bring together artifacts and communities in a common ecology that will sustain the creation, utilisation and retention of intellectual capital.* That definition is still valid, but it is specific to the time and place of its construction.

Too many technology based knowledge management systems are attempting the cognitive equivalent of raising an iceberg above the surface of ocean. It may be possible, but the energy

As knowledge management has grown and we have all gained more experience, the need for a more sophisticated understanding of the community aspect of that definition has become evident. Knowledge can only be volunteered; it cannot be conscripted. What we are finding is that formal systems are based on a degree of conscription – either carrot or stick – because of the very nature of the reward mechanisms in modern organisations. Effective learning and exchange of knowledge take place more naturally and with greater effect in the informal communities that exist in all organisations (Aibel and Snowden, 1998).

Any ecological management programme needs to create a multifaceted model that takes into account both the formal and informal aspects of the organisation. The following model (Figure 1) is an attempt to provide an instrument through which an organisation can understand its current ecology, and plan interventions that will increase organisational resilience.

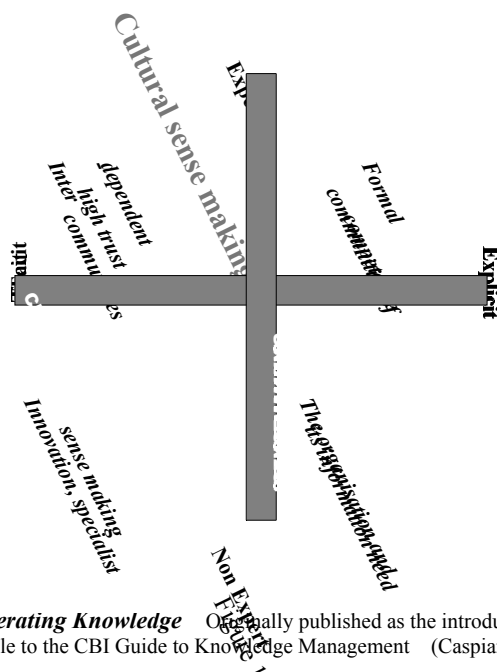
The model contrasts two dimensions as a means of understanding the different types of community that exist within the boundaries of the ecology being managed.

The dimension of culture

The two extremes of the dimension, derived from anthropology, are:

1. Tacit cultures, which define themselves by their networks, relationships and dependencies – this tends (but not exclusively) to be a characteristic of small start-up companies and the informal structures that exist within larger, more formal organisations.
2. Explicit cultures, in contrast, define themselves by their artefacts – organisation charts, documents, processes etc.

Tacit and explicit in the context of culture is not the same thing as tacit and explicit in the context of knowledge management. The Issue for many companies is that the tacit aspects of their organisation are not visible or readily definable. As a result, their management tends to be confined to the odd workshop, communication and leadership programme and a general exhortation to support company values and to be good. The problem is not just that humans are lazy and that the explicit aspects of a community are the easiest to identify and manage, although this is generally the case. The real issue is that the infrastructure and models that we use to describe organisations assume an explicit, pseudo-rational and mechanical model of decision making.



Culture in this context includes technology. Human culture does not make sense if it excludes technology – technology is the current manifestation of our tool-making capability. The issue is to see the technology

of knowledge management – search engines, document management systems, yellow pages etc. – as a tool. Good tools are those that when picked up naturally fit the hand.

The sense-making dimension

Sense making is the purpose of knowledge and the only valid objective of knowledge management. When we look at the sense-making dimension, we see two aspects or extremes that are appropriate to knowledge exchange within communities:

1. Expert language exists, which enables a group of individuals with common educational background and/or experience to rapidly interpret incoming data using frameworks developed over a common body of experience in managing the translation of that data to information. In some cases the expert language may be so well developed that it can be fully codified and contained in an artefact, such as a manual or a computer. Accounting transaction data, for example, can be converted into charts of accounts, cost variance reports etc. in an automatic manner as the knowledge of how to do so is well known and structured. Other examples of expert language include reference to a body of knowledge or thinking through academic references, the specialised use of common language (one of the big issues in integrating expert communities into the wider organisation) and the use of models and frameworks in a community instituted through common training and process reinforcement.
2. No expert language exists, either because the data–information transformation is commonplace and a universal capability within the boundaries of the organisation, or more interestingly because the circumstances are new or different. In the former case the organisation concerned may appear to have specialised or expert language to an outsider, but this knowledge is readily assimilated by new joiners through induction programmes and minimal work experience. In the latter case the newness of the situation may be immediately apparent (and possibly abhorrent) to significant proportions of the community due the radical nature of the change: a new competitor; the collapse of a stable marketplace; major regulatory change. In many cases the situation has changed but the organisation fails to recognise that change. Rigidly hierarchical organisations often suffer from this in their reliance on paternalistic figures and structures– something known to all may just not be acceptable and will be suppressed. A more pernicious example is to sacrifice long-term vision on the altar of short-term expediency. A major survey of IT companies in the early nineties on their state of preparedness for Y2K issues revealed that the second most frequent response of IT directors to the question ‘What will you do in the year 2000?’ was ‘I won’t be working for this company then.’

In both cases, the role of knowledge is to create information. This role is all too often confused by linear models showing a progression from data, through information to knowledge (and, in some of the more pretentious examples, to wisdom. Such models are not only misleading, but false. Knowledge is the means by which we make sense of data, not some meta level of information (Snowden, 1999).

Each of the quadrants of the model in Figure 1 contains a different type of community and requires a different type of support – both from technology and from management. We will now discuss each of these in turn and then look at some overall aspects of the model.

The formal organisation – non-expert/explicit

The non-expert language here is the commonplace dialogue of any member of the organisation who has been on the induction course and has a basic understanding of the role or function of the different parts of the various formal structures. Knowledge exchange tends to be explicit in nature. A large part of the ICMS requirement is to store and distribute information – the right information in the right quantity at the right time. This has become the domain of the corporate intranet, which has replaced the company handbook, notice boards and ripple-down communication.

The main reasons for using technology in this sector is to reduce cost (and the savings can be high) and improve the speed of distribution in a geographically dispersed organization. The dangers here are of over-enthusiasm and a fetish for search engines. There are some basic rules to apply:

1. The company handbook had a contents page and a usable index. It also had a common style and syntax that made it easy to use and follow. This was not done for some bureaucratic reason, but because over a long period such a format was more likely to ensure usability. These days in a large company it is increasingly common for senior staff to spend hours in increasingly desperate use of little-understood search engines over a vast collection of disconnected data (it is not informative, hence it is not information). Structure, familiar content and ease of access are key.
2. Ideally, all documents should have a ‘next revised’ date. Users are not interested in when the document was posted – assuming that it is declared as current. They want to know how long it is before that document needs to be checked for changes. They should be able to trigger electronic notification through e-mail of any significant changes.
3. Librarians are worth their weight in gold. They are human processors within a more complete understanding of the word ‘system’. As such, they are able to bring their skills of classification to bear. They are also able to help other users, guide them to search engines and sites that are easily accessed and suitable for their need. Over time, more sophisticated capabilities can be passed on and released to individual users. This is far preferable to being cast into a massive sea of websites with multiple search engines, ill-understand agents and the like. Librarianship is also a profession; it is not a resting place for secretaries or under-utilised analysts.

Communities of competence – expert/explicit

Communities of competence, communities of practice and the like are at the heart of knowledge management practice. The concept behind them is simple: the members of the group share a common expertise. This may be formal/explicit; for example, engineering graduates have access to a codified specialist language by virtue of their education. It may also be informal: all communities have common stories, rooted in their common history that can be used as a specialist or expert language. The language may refer to past failures or successes. The stories may be myths around individual hero figures.

However, in a community of competence, it is the formal aspects of an expert language that count. The reason is that the community of competence is a recognised part of the formal organisation. Membership is known and granted by virtue of position or status. In such communities, it is important to realise the following:

1. There must be controls on membership. The strength of the community is its ability to maintain a high level of abstraction in its dialogue. The higher the level of abstraction the lower the cost, and the more economical the codification (Boisot, 1998). If members are present who are not able to operate at that level of abstraction, they damage the community as whole (through inane questions or excessive posting of meaningless comments). In the same way, all documents and other material posted needs to be subject to a similar level of control if experts' time is not to be wasted.
2. Such communities can be made to self-organise. Take for instance a body of experts in a particular discipline. They need access to a broad range of academic papers, reports etc. There will always be more material published than they can possibly read, so how do they select? One way is to store the documents in a single database and not allow anyone who accesses a document to exit without pressing a radio button – say a scale of five, ranging from 'I looked at this but didn't read it' to 'everyone should read this and live their lives in accordance with its dictats'. This information can then be used to present a list of documents to other users, sorted by symbols representing the number of people who have accessed the document and the average scores. However, we can go beyond this. One of the principles of managing an ecology is to make micro interventions designed to produce a macro effect. It is also important, where possible, to construct that intervention in such a way as to allow to ecology to reveal something of itself to you. Imagine that in the above example we extend the system to allow members of the community to create lists of people they respect. We can now present them with the list of documents sorted by the two symbols for each of those lists. This tool relates to the natural way of working – I want to read what the people I respect are reading. It also enables the ecology to reveal something about itself – the lists of who respects who.
3. There are two very pragmatic issues that relate to this type of community. It must be able to replicate the virtual space in which the material is stored. A professional community wants to work on planes and trains; it does not want to be permanently hooked on to an intranet site. It is also important to ensure that there are chat areas in which individuals can exchange ideas about the more formal material used within the competence.

However, while useful, the community of competence is not a universal solution and it will have its problems. The great strength is that the community ensures that it controls entry of people and documents to the group so that they can operate at the appropriate level of abstraction. That very strength becomes a weakness, in that such groups inevitably ossify and start to stifle innovation through blackballing new members and documents challenging the status quo. Some mitigation is possible. Making sure that the gatekeeper function is rotated and subject to random audit and challenge will extend the life – but it will not remove the problem. It is therefore important to destabilise and disrupt such communities on a cyclical basis – and to put detection systems in place to trigger such destabilisation.

Inter-dependent communities – expert/tacit

Our previous two community types have been formal. We now move to the informal or shadow organisation in which trust and relationships built up over years, or in the face of shared adversity, dominate. These types of communities are the means by which an organisation is held together. Inter-dependent communities have the lowest level of visibility, but the highest level of natural knowledge exchange of any community (Aibel and Snowden, 1998). The key aspect of inter-dependent groups is that trust and dependency are at their core. This means that mistakes and errors are more likely to be admitted and discussed within such a private network, rather than a formal competence. In that case, the only mistakes admitted would be those that prove just how clever the mistake originator was!

Let us take a hypothetical example of a trusted community working in a large organisation. A project team has to be assembled. The leader of the project phones up an old friend who owns resources: 'I need a six-person team for a project that's a bit like client A, it's got the client from hell like we had in project B, I can't afford the sort of team difficulties that we had in project C.' In that very short conversation, by reference to an expert language based on shared experience the project leader has provided his friend with a far more comprehensive briefing than would have been possible through a skills profile – it has also taken a lot less valuable time.

Now what does his friend do? He identifies people he thinks will fit the bill. He phones them up and gets their agreement to take part. He then constructs a skill request to match their CVs and makes sure that the skill request goes into the scheduling system on the same day as each prospective member says they are available for work. Lo and behold, the correct team is allocated by the system. Six months later, the project is over and the leader receives a request from the manager in charge of the scheduling system to indicate his satisfaction levels – did he get the team he wanted? He did, so he says yes. The manager then uses this material to justify further investment in a more sophisticated scheduling system that my project manager and his friend will now have to work even harder to 'fix'. We have to have scheduling systems, but can't we find a better way to link the informal and formal communities together?

It is this link between the two that is important. We have to make the informal organisation disclose itself in some way to the formal so that we can use its strengths. However, if we try to legislate it will go underground. Let's look at two ways – derived from the author's own experience in IBM Global Services – in which this might be done.

1. IBM allows any individual or group of individuals to set up a private collaborate environment, subject to a token cross-charge to cover costs. There are tens of thousands of these team rooms (a Lotus product). For the past three years the author has owned one of these, admission to which is limited to those who are trusted. That team room contains working material, errors, mistakes and learning. It has been reformed at least once when a member betrayed that trust and was excluded without really being aware of it – the old team room was left in place while a new one was created in parallel with a more limited membership. That learning could not sensibly take place in a formal community of competence, one that contains rivals and the like. One of the things that my team has been working on is the use of story telling for tacit knowledge capture and distribution. The fact that this work is going on is flagged in the formal competence area (expert/explicit). This results in an early trickle of e-mails asked for details, which are answered. But story telling

becomes more fashionable; the e-mail column increases to painful levels. So I write a document answering the most frequently answered questions. I codify at a point where the socialisation pressure of the ecology forces me to volunteer my knowledge. I also know – through the questions – at what level of abstraction I need to codify. Had this only been allowed in the formal competence grouping, then either the document would have been academic, cautious and the like, or the work would have happened on a private website, or in meetings or via one-to-one communication with attachments over the e-mail.

2. A second idea – not yet an example – assumes the same principle adopted on a wider scale. But here the right is maintained by the organisation to use search engines on the private database – not to display the material, but to use key word and context searches to identify those team rooms that may be working on areas of use to tackle a specific opportunity or problem. If one is identified then the material is not disclosed, but an e-mail is sent to the team room owner describing the problem and asking if they have anything to help.

Forced interdependence – tacit/non-expert

This is the domain where we make sense of the senseless. New circumstances have arisen and our expert language developed for a different time and place. We now have two dangers (at least). One is that we may not realise that life has changed. The other is that even if some people realise it, the experts in the top right-hand quadrant will use their expertise to erect cognitive filters, to prevent the new concepts and values challenging years of accumulated specialism.

There are various ways of managing this. First, I may find that some of my upper-left communities have already starting dealing with the issue. The more I encourage experiment the more likely this is to have happened – and the faster my organisation learns.


Another is to identify those individuals with rich knowledge networks: the people who, if they do not know themselves, will know someone who does. Network analysis tools may find this out, fed with the results of successive ‘who would you ask if you wanted to know about...?’ It may also be necessary to go outside the organisation to obtain the necessary resource. The key in sense making in these circumstances is to assemble and abrade different knowledge assets. By selecting individuals who have rich private networks, you are selecting individuals who by their nature have access to such assets. All you need to do is add in time constraints and adequate resource and you have a hothouse of creativity that can innovate and compete in conditions of uncertainty.

Movements and barriers

It can be seen that it is possible to create an ecology in which each area of the model thrives. However, one is not managing knowledge if there is no exchange between the quadrants. The primary role of management in such a knowledge ecology is to manage the border areas between quadrants. The richest opportunity area is that between the two expert language groupings at the top of the model.

If we look at our model again, we can start to see some of the flows that arise.

As we make sense of the senseless, so a group of inter-dependent individuals who have pioneered that understanding develop an expert language that is as yet not fully articulated. It is very tacit, based on the common experience of sense making. Over time – and with the correct socialisation pressure – this will be codified at an appropriate level of abstraction and pass into a broader community of trainable experts. Eventually it will become commonplace. As can be seen from the diagram, the borders between formal and informal actually cross each area. Informality creeps into competence groups, for example. Equally, in the main, it takes a corporate intervention to make sense of the senseless, so the bulk of that domain is formal.



However, let us return to one of the issues in our expert communities of competence: they ossify. If we can disrupt them on a cyclical basis by moving them into unknown space, then we can take them through a process of rebirth that will render our whole organisation more agile. What we cannot do is achieve a hub position. This is the idealistic position adopted by too many members of the mechanical school. They assume that their method or software

package will produce behavior that is universally shared between experts who trust each other fully and are able to innovate in conditions of uncertainty. Such a thing is not possible; no organisation operates at the hub. However, if we increase the flow rates across the boundaries and move down the borderline between formal and informal, then we approach the hub, even if we do not reach it.

A new definition

In the introductory chapter to the last CBI handbook, I finished with a definition. In the light of the increasing differences between organic and mechanistic thinking, I would like to propose a revision.

Knowledge management is the developing body of methods, tools, techniques and values, through which organisations can acquire, develop, measure, distribute and provide a return on their intellectual assets. It is fundamentally about creating self-sustaining ecologies in which communities and their artifacts can organically respond to, and confidently proact with, an increasingly uncertain environment.

This article forms the introductory chapter to the 1999 CBI Handbook on Knowledge Management. Copies can be obtained on application to the author who can be contacted via e-mail on snowded@uk.ibm.com (please note it is snowded, not snowden)

A more substantial development of the culture sense making model elaborated in this article will be published as a chapter “Cynefin: a sense of time and place the social ecology of knowledge

management” in *Knowledge Horizons: the present and promise of Knowledge Management* to be published in October 2000 by Butterworth.

For those interested Cynefin pronounced Cun-ev-in is a welsh word which has no direct meaning in English. It describes the relationship that you have with the place of your birth and of your upbringing, the environment in which you live and to which you are naturally acclimatised. The book chapter uses the word and concepts behind it to develop Nonaka’s concept of BA.

References

Aibel, J and Snowden, D “Knowledge Management and its relationship to Organisational Change” *Focus on Change Management* July 1998

M Boisot *Knowledge Assets* Oxford 1998

T Burns and G Stalker *The Management of Innovation* Tavistock Publications 1961

T Davenport, D De Long, M Beers “Successful Knowledge Management Projects” *Sloan Management Review* Winter 1999

G Falkowski & F Foster *Organization Network Analysis: A Tool for Building a Learning Organization* Unpublished paper available from IBM

G von Krogh & J Roos *Managing Knowledge, perspectives on cooperation and competition* Sage 1996

D Snowden “A Framework for Creating a Sustainable Programme” *CBI Guide to Knowledge Management* London: Caspian Publishing/Confederation of British Industry 1998 (Republished in *Knowledge Management Year Book 1999* Butterworth April 1999).

D Snowden “The Principles and Practice of Knowledge Disclosure - first steps in an organic approach to managing in a new age of uncertainty” *Intelligence Publishing* September 1999

I Nonaka and H Takeuchi *The Knowledge Creating Company* Oxford 1995