

2.9 A NEW KIND OF COMPETENCE: ON AVOIDING MISTAKES IN LARGE ORGANISATIONS

Oliver Sparrow, The Challenge Network

Introduction

Complex activities need to be optimised in a multi-dimensional environment. The range of trade-offs, contending stakeholders and other constraints can make it hard to say what a wholly successful outcome would look like. Even in uncontested and established activities it is hard to find a recipe for action which works consistently. It is only natural, therefore, that organisations which have established a workable approach to their problems are reluctant to abandon these, to explore new territory or to re-think their fundamental drives and mechanisms.

There are a wide range of forces at work that can be expected to alter the operating environment for these organisations. The pace of change seems set to quicken. This will challenge the ability of complex organisations to adapt to changed circumstances with changed procedures, offers and relationships. It is, perhaps, unfortunate that many large organisations have responded to pressures on them to cut costs by reducing the amount of staff work that they are prepared to fund. The likely consequence is that they become less able to understand and adapt to the forces of change.

The imperatives of the times demand a different response. Understanding counters the natural tendency to err. By contrast, acting in ways for which the logic is either ill-defined or predicated on what used to be true is a prevalent source of error in large organisations.

The paper examines the common sources of failure in large organisations that are engaged in routine operations. It specifically excludes issues of risk and uncertainty, which are more than adequately treated in the literature, and focuses on how it is that purposeful organisations can lose – or, in some cases, fail ever to find – a way forward. The first sections are concerned with analysis, leading into some general prescriptions as to what can be done to generate or restore a common sense of direction. The application of these concepts to one-off, very large projects is pointed up at various points in the text, rather than treated explicitly. This paper is written from a perspective generated by experience in both large organisations and extensive consultancy to both states and very large companies. The reader is invited take from this broad brush approach whatever in it may stimulate further thought.

On the inevitability of a new approach to the management of large organisations

Failure is the normal run of things, and success is exceptional. Less than one chemical product in a hundred thousand that emerges from a pharmaceutical screen makes it to commercial success. Less than two percent of new fast moving consumer items remain stocked on supermarket shelves three months after their launch. Only three out of a hundred start-up companies survive the first three years of trading.

Most of the ideas that are passed for formal evaluation within the corporate environment will fail. All but a few percent of projects that receive third party venture

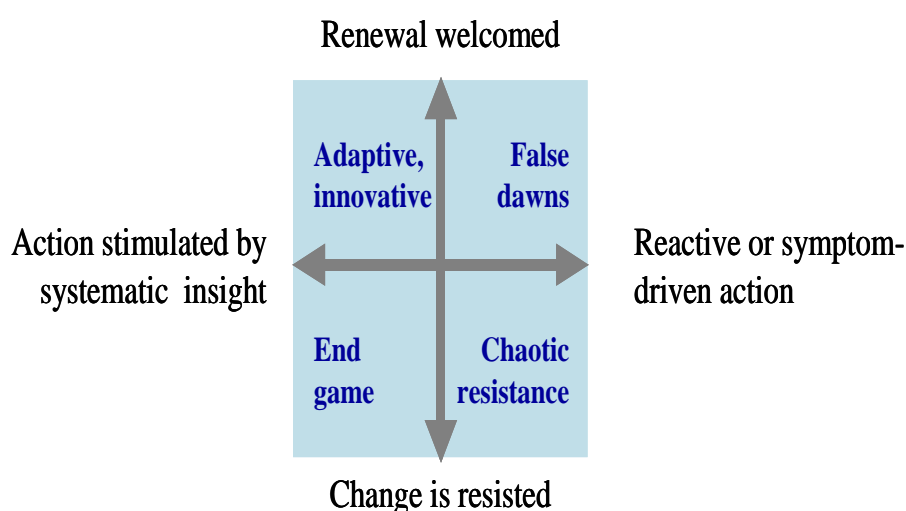
funds also fail. The commercial yields of those few projects that do succeed, in the sense of being taken to an operational conclusion, are seldom encouraging. Everyone with a background in commerce will be familiar with the wounded project, barely breaking even, in desperate need of the *coup de grace*. Persistence does have its virtues, however, as does experience. If any new activity has a 50 per cent chance of failure in any one step in a project – a gross under-estimate – and if it requires five such steps, then it is simple to calculate its overall chance of success: about 3 per cent.

Large organisations operate at an advantage in this regard. Being established, they find many of the steps already in place. The one or two new things which they have to get right are, therefore, subject to a much lower compound risk. Equally, their resilience allows them to try and try again until they succeed. Proxies for this, such as incubator parks, are able to generate extremely high success rates for start-ups, with three quarters or more of the ideas delivering commercial success. Highly focused venture houses can get away with discount rates in the 12–15 per cent region, whilst generic lenders need 30–50 per cent internal rates of return to compensate for high failure rates.

Insight counters failure. Machinery that generates insight should be a valued part of any economy. However, it is not at all clear what the tools are that create, harness or propagate insight. It is also not at all clear how we recognise and correct or quash failure in its early stages.

Figure 1 shows a traditional two dimensional matrix. The vertical axis reflects whether change is accepted or resisted. The horizontal dimensions ask whether action flows from insight or from reaction. This defines a space with four common outcomes. Most would agree that major errors are encountered most frequent on the right side of the figure. The lower left quadrant hosts an operational style which can be highly resistant to changed practice. It is a perfectly valid place to operate when the backdrop and the task are unchanging, but in the real world it needs to be informed by critique and ideas which comes from, for the most part, the upper left of the figure. The focus of this paper is upon the ways in which an organisation can be position a significant fraction of what it does – or, indeed, the totality of it – within this fertile quadrant.

Figure 1: A traditional two dimensional matrix



On the challenges innate to the knowledge economy

The effective habits of an organisation will tend to last indefinitely unless something changes. Adaptation requires recognition of change, a definition of values – of what would be desirable and unfortunate – and the adoption of measures which adapt appropriately. Naturally, the more the operating environment is disturbed or unclear, the more often such challenges will present themselves. Highly volatile or swiftly changing environments require correspondingly powerful patterns of insight if the level of adaptability is to remain constant.

This section reviews some of the forces that will lead to intense erosion of the familiar, and force active systems of scrutiny on organisations which are to adapt to this. Those which lack this tool kit, and which are consequently poor at adaptation, will survive more by chance than design. There are three such forces that are particularly pertinent¹ to the “knowledge economy”. Each of these places its particular pressures on management. Each of them demands new skills and a distinct approach to governance. The three issues are these:

- The growing importance of grappling with so-called “intangibles”.
- The importance of renewal in a world of much enhanced potential and much greater competition.
- The human resource issues which must be resolved if a coherent response is to be made to these issues.

This last concern will be amplified and developed in the closing section of this paper.

On the importance of intangibles

Economists measure economic activity in terms of ‘value added’, which is the difference between the cost of the inputs to a process and the income derived from the sale of its output. In these terms, around half of all of the value added in the OECD economies comes from ‘intangibles’ from patterns of order that are implicit in the final outcome. Indeed, where there is a physical entity – a city street, a vehicle – this will embody the outcome of countless intangible influences, from design skills to the law. It took around three years to build the Sizewell C nuclear reactor in Britain, but nearly a decade to negotiate permission to do so. London was burned flat in 1666. King Charles II devised elaborate plans to rebuild the city on less medieval lines. However, it proved impossible to proceed until it was agreed to reconstruct the city following previous property boundaries. Re-defining ownership was much harder to achieve than replacing timber, bricks and mortar. That is, intangibles often resist erosion, competition and change far more strongly than physical capabilities.

Economists speak of “factors of production”. These are inputs to the productive process, such as capital, labour and the like. A certain proportion of economic growth can be attributed to the addition of more factors – these days, typically more capital – to

¹ There are, of course, a myriad of other issues, such as demographics, and resource balances. Their impacts are generally better-understood, or anyway more self-evident – that the three pervasive and intangible forces which are discussed in the body of the paper.

existing activities. However, most activities are not additive in this way. There was a famous demonstration by Adam Smith that specialisation allowed for much more productive activity than did simple addition. (His study of a Scottish pin factory now ornaments British bank notes.) Increases in efficiency that come from ingenuity, specialisation and the like allow us to do ‘more with less’: to allow the industrial nations to use about half as much energy per unit of value added that we did half a century ago, to construct buildings with much less materials, to achieve more throughput per person, per unit of capital, per acre of land than hitherto.

Economists call this increase in efficiency “total factor productivity”. Organisations (or nations) which grow their total factor productivity are doing more with less, whilst those which do not do so are static, expanding only in line with increases in tangible inputs. Nations such as Japan showed a surge in total factor productivity in the 1960s, a rate which has now fallen away. Taiwan, Korea and now China have all shown successive surges in this measure. It often leads extraordinary economic performance. It is possible to examine the sources of total factor productivity. The case of the US shows us that recent performance is driven by better trained or more effective workers, the widespread use of information technology and the like. However, there is a large residual of when even these factors have been taken into account. This can be shown to come from insight and consequent purposeful action.

The consultancy McKinsey has published a study of manufacturing industry, based on extensive surveys in the US, Germany and India. The outcome of this was that something between a third and two fifths of the life-time cost of manufacturing came, as expected, from physical operations – that is, from cutting, welding and the like. Conventional business expenses, such as financing and managerial oversight accounted for an additional 15–20 per cent. That left 40–50 per cent of the lifetime costs of manufacturing in this huge, trans-national sample undefined. McKinsey attributed this to “the cost of gaining insight”: to managing and harnessing the intangibles, as with gaining permission to build a nuclear power plant, or to rebuild a fire-devastated London. It turns out that most of these costs were incurred in the early phase of a project, before it had settled into routine activity. They were associated with errors, delays and false starts. That is, innovation has an enormous potential cost if we manage it poorly, and – as we saw in the introduction – a very high rate of failure, redeemed only by the immense value of the occasional success.

The key point to take from this is that there is a blind spot in most organisations. Competition means that commercial organisations which are to survive must pound inefficiencies out of their “tangible” operations. However, that addresses only half of the cost structure in manufacturing, and without doubt much less than half in government, in services and in the non-profit sector, such as science. We do not have a “technology” of renewal and innovation, and we rely upon the same intuitive skill sets enjoyed by people who lived a hundred years ago.

On the central role of renewal

The life cycle that characterise all industries and much state policy are now shorter than hitherto. The reason for this is both plain and multi-factorial. Customers and shareholders have access to more information than hitherto, to a much wider range of choices and to lower switching costs. Companies, therefore, face a world in which the “winner takes all”. One can see this phenomenon most clearly in industries where

products are ranked by third parties, as with pharmaceuticals or agricultural seed breeding. Doctors (or farmers) will tend to use whatever comes out on top of the list, and none will choose anything that is low-ranked. Winner takes all².

There is, therefore, little space for second-tier organisations which are not in some way niche players or in other ways sheltered from full competition. Each major European country once played host to several car companies. Now, Europe as a whole has about the same number operating as once were found in each nation. It is widely expected that the world will have less than ten major car manufacturers in a generation. Plainly, each will become more like the other, each will pursue and increasingly homogenous customer base, each will comply with stringent international norms for safety, emissions and the like. In addition, a host of tools have come into existence through which to measure the performance of business processes and compare these to international best practice. Equipment vendors, state regulators and others add their voices to the chorus demanding standardisation around a (moving, improving) best practice metric.

This force has a surprising consequence. Just like organisms, firms prosper when they have a “niche”, a situation to which they are uniquely well-adapted. Competitors strive to eject rivals from these niches, or to re-define them. In addition, general fast change will, of course, change these niches and erode the barriers which protect them. Governments, in the shape of regulators and ant-trust lawyers, seek to increase competition, thereby directly or indirectly further reducing these barriers. This throws companies into ever-more heated pursuit of best practice, to the clear benefit of customers but to the probable detriment of employees and shareholders. Each company in a given sector becomes more and more like its peers. Their cost structures tend to come into line in all but a few industries, such as mining and quarrying. Even in these, however, new assets are pricing themselves into commonality.

This process is termed “commoditisation”, implying that each firm (or country, perhaps) becomes indistinguishable from its rivals. Each has similar capabilities, costs, market identity; and customers are indifferent as to which of their goods they buy. In economic terms, the industry’s marginal cost curve is nearly flat. As the price in the market is set by the marginal producer – the company or supply chain that just covers its costs – then this price will fall to near every player’s costs. That is, nobody makes much profit.

There are two responses to this: **first**, to cut costs and to seek scale. This was the dominant behaviour of the 1990s, with mergers and acquisitions running at unprecedented levels. Cost cutting imperatives combined with the opening of world markets, such that low cost sources of production were integrated into “value chains”. In these, goods often pass through and fan in from dozens of producers before they reach the ultimate customer. Japan was, of course, the innovator in this field, outsourcing the energy – and later, manpower-intensive aspects of its industry to, in particular, Korea and Taiwan. The consequence has been control over these activities has shifted out of Japanese hands, leaving the more cerebral aspects of design and organisation with the outsourcing companies.

² Indeed, such industries are forced to seek scale. For example, research activities will produce more products if more is spent on them. If the quality of research remains the same, therefore, a large activity will “take all” disproportionately more often than a small one. The continual drive to merge or acquire in such industries follows from this.

The resonance of this with the issues of the knowledge economy, of creating and managing intangibles are unmistakable. The **second** response to commoditisation is, therefore, that of renewal and innovation. In part, the organisation needs to find new sources of added value to replace those which have been commoditised. In part, however, it needs to restrict itself to those aspects of existing commodity activities which others can not do for themselves. Most of these aspects are either intangible – relying on informal networks of people or relying upon more defined or inalienable legal framework – or they are extremely complex, such that commodity suppliers cannot economically replicate them; or both of these. Activities of this sort will, however, inevitably succumb to erosion into commoditisation over time: patents cease, contracts become irrelevant, technologies alter and so forth. Innovation, finding new things to do, new ways to do established things or new frameworks and places in which to do them is the only sure bulwark against this.

There is, as yet, no defined ‘science’ of innovation. Rather than spend resource on risky activities, most firms (and policy makers) prefer to seek accretional rather than discontinuous change. The whole industry tends to make these changes in parallel and, if anything, this tends to accelerate commoditisation. In addition, many industries which rely heavily on innovation – such as the pharmaceutical industry – are also heavily regulated, always faced with litigation around new products and in other ways deterred from innovation. As a result, the output of new compounds per unit of investment has fallen steadily, despite enormous advances in the underlying science.

The forces that lead to commoditisation will continue to intensify. Nations and firms which do not take bold steps will be trapped into a cycle of cost-cutting which weakens their capacity to make just such leaps. Those which leap without capability or insight will find the fire as hot as the frying pan. The forces that drive change are multi-dimensional. We have looked at the internal dynamics of commerce. However, external to this and largely independent of it, science and technology are growing in weight, quality and integration. The disruptive potential of such developments is plain. Organisations which can harness this knowledge will be able to avoid commoditisation; whilst those organisations which cannot do so will not.

It is said that more science is being done in the first decade of the Twenty First century than was accomplished throughout human history up to 1970. Certainly, the papers published in journals such as Nature and Science read, in some cases, almost as science fiction. Quantum uncertainty is to be harnessed in computers which can undertake all branches of an evaluation simultaneously. A cubic kilometre of Antarctic ice is being turned into a kind of telescope; whilst other telescopes with flexible, adaptable mirrors peer at the immense black hole in the centre of our galaxy.

The genetic code of organisms is now routinely described, including those of extinct Neanderthal humans. It may soon be commercial to sequence that of an individual patient, providing a prognostic on that person’s future health and, perhaps, social, intellectual and other potential and weakness. Clinical interventions that are based on the ability to regenerate an individual’s tissue, perhaps to the extent of growing replacement organs may greatly extend enjoyable life. The human mind is under dissection with elegant tools, identifying the location in which specific tasks are undertaken. Complimentary animal models of cognition are developing quickly. Gene manipulation allows, for example, a given class of neurons to be turned on or off with a burst of coloured light. This, in turn, allows categories of processing to be triggered remotely or inhibited at will.

Machine-based emulation of high-level cognitive tasks will be a feature of the 2020 period, with vast social and commercial implications. Middle management tasks could be automated and distributed, given utterly focused value systems. The implications for the offensive and defensive side of state and commercial security are self-evident.

On burgeoning human resource


There will be more graduates in the work place in 2020 than there were people alive in 1900. The information technology of the period will be remarkable, allowing semi-autonomous agents to act as assistants, data miners, advisors; permitting brief meetings to be convened spontaneously across continents in ways which match current physical presence; carrying the analogues of reputation and track record across organisational and national boundaries. Any opportunity will be grasped quickly and acquire the scrutiny of experts from around the world; any fixed position will be exposed to continual scrutiny, erosion, arbitrage.

Intangibles are held in three ways: in formal agreements and processes, in informal social behaviour and in individual heads. All experience to date suggests that innovation and renewal comes chiefly from the second and third of these. That is, for organisations to avoid errors and to achieve success with new initiatives, it is the behaviour of the people who are involved which is of paramount importance. Why such behaviour is often counter-productive is the subject of the next section. How collective behaviour can be made more useful is the subject of the closing chapter.

On the competent organisation and its enemies

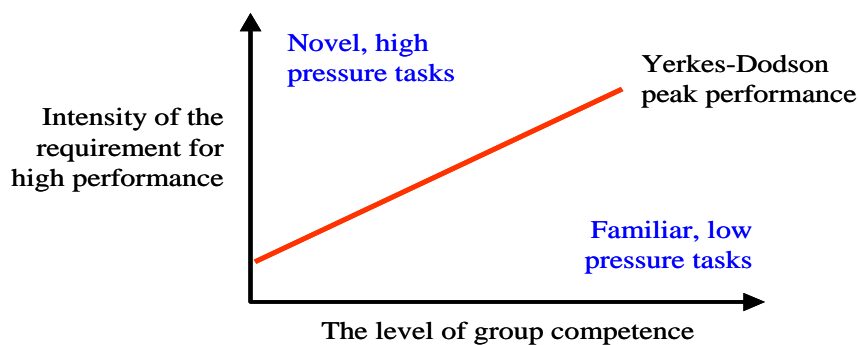
Psychologists who are concerned with the performance of people at work have long drawn on the views put forward by Yerkes and Dodson in 1908. People respond to a challenge in ways which reflect the intensity with which it is presented. At very low levels of intensity, arousal is also low and performance is weak; as the intensity increases so too, arousal rises, first to match it and later to be overwhelmed by it. The Figure 2 shows this schematically. The downward trend that occurs beyond the peak occurs because people are overwhelmed either by the sheer volume of a repetitive task or the cognitive demands of a more complex set of activities. Either way, stress increases beyond the optimum point, and people are easily distracted, irritated and inclined to make mistakes.

Figure 2: Performance and levels of arousal/pressure



This observation – which has been demonstrated in thousands of trials during the Twentieth century – can offer some useful lessons. We can develop these in the related Figure 3, which is shown below. This shows the requirement for intense performance as the vertical axis, and adds a new dimension – the capacity to act effectively – as the horizontal axis. The onset of Yerkes-Dodson peak performance is shown as a line that slopes upwards in response to increased collective competence. Activities above the line push the organisation into a sub-optimal working regime, and the rate of error-making will increase. The term “competence” has a special meaning for psychologists. It refers to the sense that groups and individuals have that they are operating in familiar territory, are empowered to act and are otherwise ready to respond.

Figure 3: Onset of Yerkes-Dodson peak performance



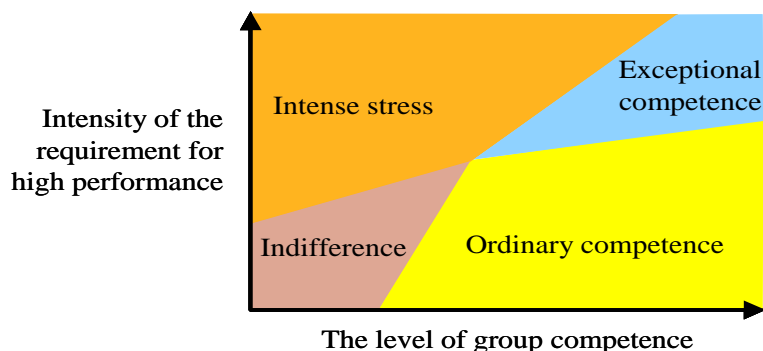
To understand what is meant by competence, one might imagine the state of mind of a group of young hunters setting out on a fine spring morning at the end of the last ice age, surrounded by plentiful game. Not only do they know what is expected from them, they know that they can deliver against these expectations. They know their peers intimately, their strengths and their skills. They know that what they are doing has general approval, and that there is an element friendly competition within the band.

The term competence therefore bundles trust, legitimacy, capability and confidence. It is a fine state of mind, and one that can be held by people who are, under an informed eye, either correct or completely wrong in their view of the world. It is no guarantee of a sensible stance towards events, only an assurance of confidence. Taking this caveat as understood, however, organisations which work to increase their competence also extend their ability to deal with otherwise stressful situations. They shift the gradient of Yerkes-Dodson threshold upwards on Figure 3. Training, proper procedures, good team interaction and other factors all contribute to this aim.

The reader may feel that this analysis re-states the obvious. However, Figure 4 shows us how this way of looking at the situation allows us to map the terrain in which the organisation is likely to have to operate. The space has been divided into four regions, in which workers feel general indifference, intense stress, normal everyday competence and, in the darkest shade, are propelled into a state of exceptional competence. The movement between these zones is telling. For example, a group with low innate competence will oscillate only between the brown zone of indifference and the orange area of extreme stress. People who are trapped in such an environment make

huge and frequent mistakes. They often work exceptionally hard, but without engaging with useful machinery.

Figure 4: Shifting the gradient of Yerkes-Dodson threshold upwards



Somewhat higher up the scale, competent people are propelled by unfamiliar challenges into indifference and, ultimately, into stress. The indifference comes with a mental caption that this is “not for me”, or that it is “not our business”. The stress arises when the responsibility to cope arrives without corresponding legitimacy, capability or insight. Highly competent organisations have extremely elastic boundaries to what they are prepared to think of as their proper area of concern. They think about the system in which they operate and about how that system works, how it impinges on them and they upon it. They are aware of change within the system and the implications which this has for adaptive responses from them. When the demands of the situation increase, therefore, such groups make the transition between ordinary and extraordinary competence. Naturally, there are limits to this and even the most extraordinary team will hit its limits and lapse into stress; but this happens under far more extreme tensions than occurs at lower levels of competence.

One can also turn this observation on its head. That is, teams which are competent hit their Yerkes-Dodson optimum at far higher levels of throughput, in far more novel situations than do the less competent. Competent parts of an organisation become – or are labelled as – the elite. This has some unhappy consequences.

Parts of most large organisations have usually been wholly bureaucratised. This may have happened by chance, or as a result of their being mummified by well-meaning total quality management systems. This division of labour usually works very well in the short term, provided that two consequences are fully accepted:

- First, the bureaucratised parts of the organisation remain safely sequestered from the wider world, and the organisation is not attempting to change itself in ways that affect them. This is a vain dream rather than a reality for most situations, of course.
- Second, that the rest of the organisation can afford to cut itself off from the insight, knowledge and social connections that may be vested in the bureaucratised parts of it. It cannot, of course, to allow the people in the bureaucratised world to take a substantive part in its decision-taking, because their views are necessarily constrained and ‘incompetent’, in the sense used above.

These conditions are, of course, usually an unfortunate state of affairs for the long term health of the organisation. We will explore why this is so in the section which follows. In addition, however, an organisation which has fallen into in this form carries within itself a group prone to hysteria and sudden breakdown when conditions change or become intense. Many senior managers will have had experiences which confirm the existence of their forgotten mine fields, that appear tranquil pastures until life becomes difficult.

Unhappily, one of the dominant trends in business culture over the past two decades has been to extend such bureaucratisation. This is often done in the name of quality management and, formerly, re-engineering; whereby every activity and process, situation and event is codified, allocated to individuals and groups and turned into a giant algorithm that runs on a machine made of people. This is done because such structures work well with IT systems, and because they define cost centres and thus help to cut costs. In addition, it aligns perfectly with some of the more malign but apparently objective HR practices, such as metric-related performance-based pay.

Traditional organisations were set up on military lines, with an officer class of the competent and a soldiery of those who did what they were told. The organisation was split into division with different specialities. These interacted only through their officer class. Some further interaction occurred as a result of sub-contract to service activities such as IT or HR. This support was, however, largely void of intellectual input: that is, IT were there to plug in computers, not to suggest how the division might rethink its mission.

The inadequacy of this became apparent between 1970–80, with the popularity of matrix organisations and the expansion of central planning groups. The trend was reversed in the 1990s, both with the cost-awareness that stemmed from shareholder value movement, and also driven by the need to change organisations to be more like machines so as to harvest what the IT systems of the time could offer. We have already mentioned some of the features of this approach.

It has therefore become commonplace for an organisation to revert to the historical position of having a small but competent cadre based upon a wide range of bureaucratised worker bees. As the bureaucratisation provides an exact description of what is supposed to pass in an out of each cell in this beehive, so individual cells can be outsourced, sub-contracted and rendered even less accessible to the organisation. Each operates in its microcosm, separated from the machinery that generates general competence. We have noted some of the unhappy consequences of this.

On how complicated organisations can attain general competence

Organisations are often wrongly compared to organisms. There is a sense in which the analogy can be useful. Organisations do take in resources and excrete wastes; they do yield useful products; they do reproduce and die; they collaborate and they fight; they sense information, take decisions based upon it, turn those decisions into action. However – and it is a very important exception, given what we have so far discussed – this analogy reduces most of the staff in an organisation to the status of mere agents of the will of the all-encompassing eyes, ears, brain. They are incompetent, in the sense used in the preceding section. The organisation acts as though all potentially useful ideas, discussion and direction are derived from a small, closed cadre.

Complex organisations cannot possibly process all of the skills that they contain, or all of the insight that they access, through a relatively tiny group of senior managers. Senior staff is often decades away from their training, and their technical insight into important areas is obsolete or non-existent. Their connection with potential collaborators, customers, regulators and others is increasingly stylised and schematic. That said there are deeper reasons why this analogy is misleading. There are two such reasons, and together, they give us a helpful point of departure in thinking why companies make mistakes:

- First, organisms have built in drives – in higher organisms, emotions – on which all of their actions are predicated. Theorists may pretend that organisations are the same, but this is simply not the case. There is no unified set of felt values in a corporate entity, only the filtered and distorted sum of individual responses.
- Second, the nervous system of higher organisms generates a synopsis, an “organism-wide” view of the current situation. There is nothing equivalent to this in human organisations that arise spontaneously. Indeed, immense amounts of work are expended to create the shadow of such a thing and to keep it up to date. The pursuit of such “sensor fusion” is of great importance to the military, where information that is gained from many sources needs to be fused into just such a synopsis. Research on this is still a long way from achieving what a mouse appears to deliver without effort.

Like a leaf on a stream, something that has no capacity to adapt to its environment can, of course, wander into situations which are destructive to it. One cannot call this a ‘mistake’, however, for things which make mistakes have, by implication, the ability to take informed decisions. An informed decision consists, in principle, of three distinct and largely separate modules of activity:

- **Model creation:** forming an understanding the operating environment, and being able to interpret events in it.
- **Value definition:** deciding on what is desirable and to be avoided; on setting balances between these variables.
- **Option identification:** deploying the model in the light of the values and of current information in order to define what is possible; to set priorities amongst these options and to take the consequent action that is needed. Generic options – the “kind of thing we’re looking for” – are often captured in the **narrative** of an organisation. This is the story that it tells itself about its situation and its goals, often in very indirect ways that are bound up in war stories, metaphor and social or operational style. We discuss the narrative in depth somewhat later in this paper. The narrative carries much that we have already discussed as shared organisational competence.

Aware organisms appear to handle this bundle of activities as a seamless whole. Complex organisations are able to do so only with great effort. However, those that possess an informed narrative about “who we are and what we are about” have generated something which is truly valuable. It allows then to innovate, to define the concrete steps forward that we often call ‘strategy’, to explore the less clear options with a view to broadening their insight, to develop people, to draw boundaries to their conduct, to their range of activities, to their aspirations. It permits them to navigate

within uncertainty, not through any illusion of possessing certainty but through a sinewy, reflexive adaptiveness that comes from shared insight and reduced ambiguity about goals, conduct, tools and aspirations. Like the hunting band that we discussed earlier, they feel competent, legitimate and certain of themselves.

The concept of the 'narrative' may need some amplification. As the term is used in sociology, it refers to the way that people in a group discuss and are given permission to debate identity and aspiration. As such, it has no quality parameters whatever that are associated with it: it can be a well-calibrated and informed analysis, or a patchwork of folk sayings, imposed dogma and the like. It can be hard for a group to criticise its own narrative, not least as many of the terms of reference of such a critique are embedded in the narrative itself. Denial, blaming the messenger or the sources of change is all too common responses which led to eventual crisis.

One clear quality parameter for a narrative is how well the guidance that it offers meets the aspirations that the group has set for itself. This is most likely to happen when the current insight that the organisation feels that it has is subject to more or less formal and informal critique, from as many angles as possible, involving as many people as possible. This implies that such critique must be embedded in some form of enabling social process. This is no more than the legitimisation of the idea of debate – as opposed to command and control – and the provision of a listening ear. It is also the broad accessibility of ideas and thinking tools that many organisations may tend to reserve for the elite. That is, the interpretive model, the values and the possible options open to the organisation need to be widely disseminated, of course tuned to the preoccupations of different parts of the organisation.

In part as an antidote to the preoccupation with operational matters, and in part to enable formal planning ideas to be introduced to the debate, internal conversations need to be supplemented by calls for comment that interact with the ticking of the grand clocks in the organisation. Business units, for example, should be required to take in the key issues – environmental concerns, perhaps – and formally reflect these in the plans which they bring forward if they want those plans to be approved. The organisation and the individuals that comprise it are thereby informed about the operational engine in which they are embedded. They can access accurate and calibrated views on how 'it all works' and, therefore, will be able to focus on what matters the most. A view of 'where we may fit into this' arises naturally from this. It leads to organic debate on the actual and self-imposed limits and imperatives that have been adopted, or which should be adopted. What the organisation needs to know in order to take better decisions becomes a self-defining agenda.

As must by now be clear, a knowledge-based organisation or a democratic state has to engross as much expertise (and implicit permission) into the narrative as possible. It has to do this without diffusion, aiming always for the sharpest focus, the minimal description that catches the overall richness. The aim of the narrative is to present a framework from which debate can depart and against which debate can react. It is not a statement of final truth. A well-developed narrative represents a connected, solid assertion of the core truths that affect, enable and bind the organisation or society. Nations, organisations and individuals which lose their narrative lose a deep ability, which is to find spontaneous order through local-level initiatives. Markets are, of course, a common example of such emergent order, as are the tacit rules by which most societies manage the vast bulk of their activity.

Coherent behaviour can arise from two sources. First, it can be mandated. Here a grand imperative can in some way be imposed – perhaps as a grand plan, perhaps through the workings of an insuperable economic force. Such structures are always present, by choice – to police exceptions and default, for example – and as a fact of life. The grand plan has a chequered history, however, not least when those directing the masses have limited insight, poor information or special interests that are not widely shared. We have already seen that this is the natural condition of senior staff in large organisations.

Second, coherence can arise spontaneously, by emergence from the local behaviour of many independent agents. This will always occur when, despite their independence, members of the group share a common set of values, insight and expectations and, of course, when the behaviour or the operating environment meets those expectations. This form of organisation is always resilient, adaptive and innovative. It can be prone to erratic or extreme swings in behaviour when insight is poorly calibrated or when debate is stifled. It can lead to individual behaviour that is collectively destructive, which is why a degree of mandate is always required to complement the ebullience, creative and erosive nature of emergent behaviour.

This is a fine balance for complex organisations to manage. Groups which lack a narrative can only gain coherence when presented with grand imperatives: that is, when their agenda is set for them. They will change only when the command elite tells them to do so, and will tend to struggle against rather than for such changes. The elite may have limited insight, and emulate what peer organisations or nations are doing rather than seek differentiation or innovation. Groups which have a narrative, by contrast, can achieve spontaneous coherence, entirely because their individual behaviour is coordinated by what they share. The elite enable such debate, and are surprised by what it uncovers. The full expertise of the organisation is deployed on the issues which the narrative defines as being critical. The organisation acquires a ‘common sense’ about what matters, what is possible and what is desirable, and finds options which answer to these in the most unexpected places.

In companies, for example, the main connections into the realm of technological change occur through the contacts between junior staff, both internally but also with the wider world. They need to know a good thing when they see it, and they need a conduit into which to place their partly-baked ideas. Much the same can be said of middle management staff and the options that present themselves about operational change: outside collaborators, suppliers, customers; novel approaches to information flows, purchasing, manufacturing, legal possibilities. Again, those presented with latent potential need a framework against which to perceive this, and machinery through which to do something with this perception when the spark has been lit.

A fine narrative is not a perfect antidote to error. People may make choices under perfect insight about the systems within which they operate, but do so subject to imperfect information. Every financial trader probably makes this kind of mistake every day of his or her working life. However, they can also make another kind of error. This stems from poor insight and from weak systems by which to generate this insight. Such structures may not exist, may exist but fail to interact with each other, or may simply be imperfectly configured for the task in hand. This kind of error – in which the organisation or the individual are bombarded with data which they cannot filter or prioritise, where there is a fundamental lack of an interpretive framework against which to work, is a very much more common source of error, block to adaptive change and

generator of what might be called corporate anxiety. Organisations which strive to be more need to leave floating down their particular stream of events and take a grip upon the three parallel and non-linear processes which have been described concerning: model creation, value definition, and the creation of a widely-shared, coherent narrative that in turn leads to options and equally coherent choices. A ‘machine made of people’ needs to have tools and processes in order to be able to do this.

One should note, however, that it is not possible to arrive at a formal model of the ideal information management system for *any* organisation. Activities differ very considerably. One can describe this difference in a space which is spanned – described by – many independent dimensions. Examples of these include (but are far from confined to) the following:

- The maturity of the activity.
- The degree to which it is open to bureaucratisation and automation.
- The scale and time frame.
- The number of stakeholders and the resulting spread of value systems.

Plainly, systems that work for some parts of this space will be inappropriate for others. What is advisable for one part of this space (a small, innovative start-up) will be quite wrong for a mature government department processing routine issues; or for an engineering mega-project in the middle of its implementation phase. That said, what varies between these is *how* an information processing system is instantiated, not whether there should be one. Any organisation which does not understand its operating environment, which does not have a clear value balance set and understood, which has not decided on a course of action in the light of this will succeed only by accident. It will be prone to make mistakes. Bearing in mind the caveats about scale, life stage and so forth, we are now in a position to look at each of the three functions – model creation, value definition and the creation of the narrative of clear options. We can also see how these steps most commonly fail.

Model creation

When used in the sense of this paper, a ‘model’ is a way of thinking about a discrete domain of activity. That is, it is not a way of thinking about the entirety of experience – the meaning of life, so to speak – and nor is it a detailed or even numerical description of some tiny subset of this. A good model has at least four important quality parameters:

- First, the model must be correct. Obvious though this sounds, we tend to take the core narratives about how much of our world works entirely on faith. *This* is how people are expected to get along together in this society. *That* is how we assign leadership and power in this society, company, family. Research shows that it may take considerable dissonance between what we believe and what we experience before we abandon a familiar model. What here the criteria should be for a “correct” outlook depends, of course, on the situation: a general guide to conduct, a scientific truth, a proof in law. Nevertheless, robust models are calibrated against experience and broadly predictive of events. They are self-contained, and do not have to call upon oracles and interpreters when things go wrong. They are used without intermediation (that is, to operate without apologetic interpretation along the

lines of “what the oracle meant to say was...”). Generally, resilient models are grounded in a network of related views that we may have about how other things work. That is, we may feel that a purely economic interpretation of events can usually be imported seamlessly into a socio-political question and so contribute to debate about this. We are much less happy when deeply dissonant models collide. This is not so much a question of different models that disagree as to their predictions. Rather, the very terms of reference on which they rest are not mutually compatible. A religious perspective, for example, may not knit well with the economic model that was mentioned above. This kind of dissonance is hard to resolve, easy to ignore and a frequent source of error.

- Second, the model must be accessible to those who would benefit from taking decisions in the light which it casts. Exceedingly complex or technical ways of thinking are not helpful to people who cannot access the concepts that this requires. Models which are not a widely shared place those who do access them in the role of dubious oracle, a person whose view you take on authority or trust, and to whom you cede your decision-taking powers.
- Third, a model needs to operate at the right level of abstraction. A manual for your new computer is neither useful if it is pitched at the level of the individual transistor nor if it offers an oversight on information theory, the role of computers in society and so forth. You want to know where the “on” button is to be found, and the key things that you need to understand in order to be able to use it. It is often easy to capture some (not very important) aspect of an issue in a very detailed operational plan, spread sheet to other tool; but in doing so, to ignore much more important but less tractable issues. (One is reminded of the man found looking for his lost keys immediately under the street light, not because that was where he dropped them, but because that was where he could see in order to search.) A good model, in this third sense, needs to illuminate the entire domain in which choices are to be made. In doing so, it often suggests linkages which are not immediately apparent. It will certainly explore complexity, and the good model is one which has had this pared back to the minimal level, exposing the key structural features which affect the way in which the “engine” under discussion operates.
- Fourth, the excellent model is informed by data, wherever this is possible. It relies as little as possible upon rules of thumb, grand theory and extrapolation from past experience. (Truly excellent models allow one to price uncertainty, and therefore to set a value on the research which will generate the data that will reduce or manage that risk.) Data can be “hard” or they can be experiential, anecdotal, issues of experience and intuition.

An important facet of model adjustment and calibration, upkeep and upgrade is it must be open to, and indeed attract, constructive critique. People should work to identify ways in which it can be simplified, and ways in which it needs to be extended in order to offer a better service. In general, no model is ever to be deemed perfect, and the routes to its improvement must be clear to potential contributors to the process.

The fine model is, therefore, in close alignment with the system that it is trying to describe. It is widely accessible to people who need to navigate within that system. It employs a useful level of abstraction which combines completeness with a direct interface with the quality of issue which its users face. Finally, it is data driven where possible, and always open to improvement, critique and reconfiguration. Poor models,

of course, fail to achieve one or more of these quality parameters. However, poor models can have their passionate defenders. Poor models can serve as psychic wall paper, covering the conceptual cracks in an organisation. Destructive models may, however, actively mislead and cause people to misinterpret events, to take measures which are actively harmful and to attack those who are trying to help. We shall look at some of these pathologies in a moment.

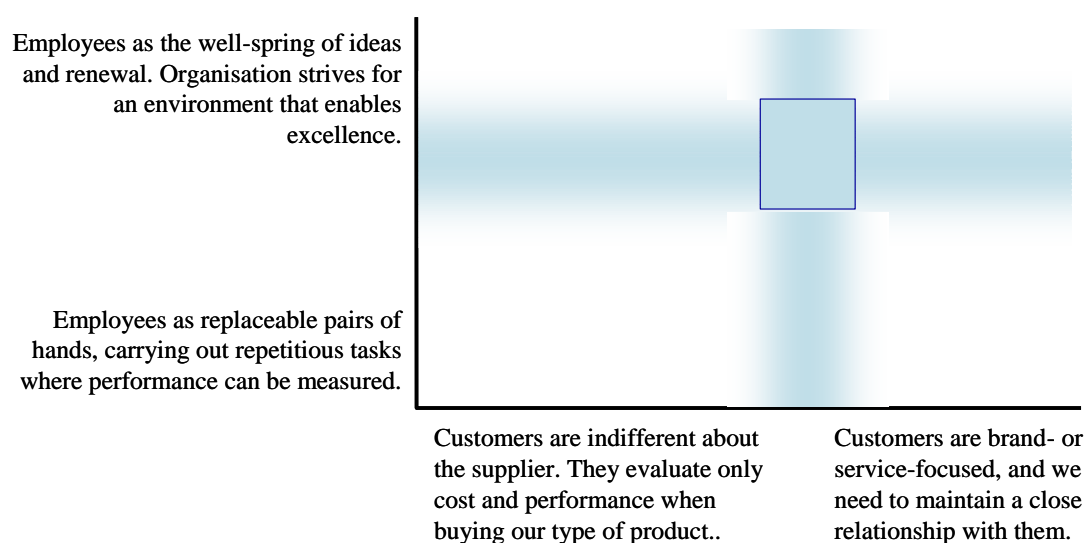
Value definition

It has become fashionable to talk about “values”, usually implying that these are something which organisations choose to possess, and through which virtue is somehow caused to abound. This is not the sense in which the word is used in this paper.

Organisations have to make choices about their operational style along many dimensions – in how they treat their staff, for example, or the position which they want to adopt for their brand in the marketplace. There may be strong reasons for adopting the position that has been chosen. For example, other parts of the brand “space” may be occupied by powerful competitors. To move into this ‘occupied’ brand space would be to attract aggressive behaviour which the firm is ill-suited to counter. Often, however, there are no firm reasons for adopting the position which has been chosen except, perhaps, echoes of the past and the collective identity or ethos that emerges from this.

Figure 5 below employs two arbitrary dimensions of this sort. The vertical axis offers a dichotomy about how employees are seen and treated: as the source of precious insight, or as replaceable pairs of hands. Either of these might meet the facts about a given activity and location. However, the organisation in question operates in a band, symbolised by the horizontal shaded blue area. The other axis considers customers as indifferent value-seekers or as brand – and producer-focused. How this organisation in fact operates is shown as a vertical shaded band below.

Figure 5: How employees and customers are seen and treated



Where these bands cross, a square shows the resulting “culture” or set of operating values. The activities with which the organisation feels comfortable map into this square. Other viable location in this “space” – in the lower left, for example, where nameless toilers make cement sacks for an indifferent world – are areas in which the people who work within the darker shaded box would feel uncomfortable. These are not their “values”. Please note that this says nothing of the efficacy, putative morality or appositeness to these values to the situation in which the organisation finds itself. It is simply a statement that its staff feels happier, more comfortable, with ways of operating which map into this part of the space. Values do, however, have quality parameters, and they exhibit these in much the same way as do models. Elaborating upon this, two points need to be made:

First, the values need to cause people to respond to their model in ways which promote the survival of the organisation (or its society, or corresponding values of an external group who feel reason to comment upon it.) There are two aspects to this. The model may prescribe behaviours which require the organisation to migrate to a new part of ‘value space’. (For example, changing market conditions may require a firm to move to the lower left corner of the diagram, in which the cement workers toil.) This may provoke the organisation to anything from internal dissent to the complete denial of the model, often shooting the messenger who suggests it. In addition to this, however, organisations may disagree with the values of more or less important stakeholders. It is usually swiftly fatal for commercial organisations to disagree with their customers, although state enterprises seem to thrive in such a situation. However, there are many lesser sources of pressure in which such disagreement can be a chronic source of friction, but not be enough of a force to shift the organisations’ values. There are illustrations in labour relations, for example, or in response to external pressure groups.

Plainly, it is desirable for an organisation which is under pressure to be able to change its values, provided that in doing so it remains faithful to its model of how everything else works. Organisations which change values without thinking about the model tend to wander into crisis. US car manufacturers mortgaged their future to worker pension commitments, for example, without thinking through what this did to their long-term competitiveness; whatever their motives, the dislocation has led to their general near-insolvency.

Second, all parts of the organisation should have adopted the *same* set of values. If different parts of an organisation inhabit different parts of the value space, then they may have difficulty in finding common ground, insofar as they cannot agree on what an acceptable outcome will look like. Many organisations manage this situation by finding forms of words which blur necessary distinctions. Their divisions operate on the basis that nobody should look too carefully at what is happening ‘across the fence’. However, organisations are necessarily heterogeneous. They contain the tender shoots of start up enterprises and the hard hands of commodity businesses that are being run for cash. Their optimum values are distinct. Values are, therefore, necessarily nested. If there is no level at which the component parts of an organisation can agree, then it is arguable that they should not be a part of a common structure.

However variable practice needs to be at the unit level, common values must lie across the organisation as a whole if it is to preserve its unity. The political concept of subsidiarity – of taking choices at the most local level possible – has many parallels with this. If a region cannot accept the binding principles that are supposed applied to all of the participants in the political union, then it really cannot remain a part of that

union. At the highest level, therefore, the organisation can set out the need for change that it sees ahead. There will be renewal and retreat, expansion in this area, retreat in that one. The deal for stakeholders (customers, employees and so forth) will fall within general parameters. In those terms, the best balance for stakeholders within a given division is thus, and in another, so. The outcome is different, but the motive (and the self-imposed limits set by values) is common to both.

Values are, therefore, statements of what matters, and where organisations choose to cluster their responses to these variables. What is thought to ‘matter’ is generally defined by the model. The responses are set by a mixture of rational (that is, model-based) calculation, and collective-subjective feel for what is appropriate; and habit. ‘Good’ value systems are internally coherent, match the model and are shared across the organisation in such a way that deviation from them is always a conscious and disciplined response to local conditions, subject to evaluation and review. In a tightly coupled litigious world ‘good’ values recognise a duty of care to stakeholders and to the system in which the organisation is embedded. They are not, however, explicitly “ethical”, either by design or outcome. It is the job of legislators to translate whatever categorical imperative a given society may feel into a framework that the organisational model can recognise, and the value system encompass.

Option selection and the organisational narrative

Narratives have already been discussed in some detail. There is, however, a further distinction to be made. Knowledge which is held by an individual or an organisation can be ‘declarative’ or ‘procedural’.

- *Declarative knowledge* can be written down, taught symbolically and transmitted impersonally.
- *Procedural knowledge* can only be acquired through experience or by emulation. You can be told how bicycles are ridden, but you can learn to ride one only by going through the experience yourself.

Narratives allow us to transmit procedural knowledge. Children who absorb tribal lore at the fireside are gathering deeply-held insight into proper behaviour, into how to think about and approach common problems. When they use this knowledge, they often do not even know that they possess it, for it has become a part of their deep responses to life and its challenges. Much the same is true of national or organisational narratives. Not only do these dictate responses to common problems – as opposed to consulting the rule book or FAQ – they do so in terms of what feels like common sense, affording the confidence to make quick, flexible responses which are nonetheless harmonious with what the rest of the organisation (society, company) wants.

This leads us to the issue of **option creation**. A well-calibrated and informed model and clear values do not of themselves create options. Rather, they can be induced to create a narrative, through which people are able to perceive potential through the fog of operational detail. Creating an option about which concrete things can be done is a complex process if another set of mistakes are not to be made. Options are presented to us when we have a narrative that is strong enough to allow us to perceive them. In the absence of values, of a model and their fusion into everyday life through the narrative, we do not see opportunities. If a chimpanzee has seen another using a stick to extract

termites from a mound, then a loose stick is a potential tool. If this option has not been grasped, however, then the stick does not offer an option: it is just a stick.

A few corporations are run by single figures that see all, understand all and command all. Generally, however, commanding individuals are exceedingly good at getting a network of people to interact appropriately, and to induce its members to communicate with each other. Virtually all large organisations run through consensus, with the members of the elite individually able to question and hinder projects, but almost none of them able (or willing) to force a project through the system that is trying to resist its adoption. Options that are to be adopted must, therefore, present themselves as self-evidently attractive both to the model (“that’s where we want to go”) and to the values (“that feels right”).

Commonplace sources of error

Organisations can continue for amazingly long periods *without* a clear model of their operating environment. Quite senior individuals can be almost completely ignorant of everything that lies outside their immediate sphere of concern. Organisations may fail to define their values, perhaps in order to avoid conflict. Some are riven with contested values. In each case, it is extremely unlikely that such an organisation will take easily to options that are placed before it.

The reasons for this are clear. There is no clear narrative, so people will lack a collective, intuitive sense that the option is what is needed. On the contrary, new possibilities are likely to be seen as the means to stir up debates that many wish to remain quiescent, or to raise questions in the heads of individuals which they do not wish to address. Conflicted organisations are almost invariably bad at innovation, bad at spontaneous adaptability to changing conditions, and bad at maintaining internal and external dialogue. The last, of course, makes it even less likely that they will improve their model. Indeed, they tend to cling to the status quo until action is forced on them by crisis. Like individual humans thrown into the same situation, they often react by attributing blame rather than solving the issue, by denial of the problem and by intensifying measures that used to work, or against which it is hard to make a case, such as cost cutting, reorganisation or financial manoeuvres, such as acquisitions.

Individual personalities have a role to play in this. There has been considerable research into the cognitive styles of people who rise in large organisations, and the types which thrive in those which have got themselves into the quandary that we have just described. For purposes of discussion, the many types that are found can be collapsed into categories, the Hedgehog and Fox personality types. People who can be classified as Hedgehogs are happiest in a closed problem domain, in which standard tools and focused effort allow them to compete with their peers. Foxes, by contrast, are at their best exploring new terrain, developing alternative strategies. Their goals are largely internal and seldom benchmarked by competition against their peers.

Groups of Hedgehogs constitute wonderful engines by which to deliver against unambiguous tasks and tight deadlines. Groups of Foxes neither enjoy such tasks nor perform well at them: indeed, Foxes tend to operate in flexible networks rather than closed hunting bands. However, Hedgehogs are often baffled, and perform at worse than chance, when asked to extend the borders of their current activities, to predict events and to mitigate new sources of risk. Foxes excel at such tasks.

Hedgehog types easily drop into over-confidence and over-focus. A complex policy problem is “nothing but” poorly functioning markets, imperfect information, elitism or whatever is the fashion of the moment. The solution to corporate woes are “nothing but” cost cutting, key performance indicators, risk analysis. The cerebral approach that leads to isolated “mini-models” is extremely seductive to Hedgehogs. The bold reduction of a problem to an over-simplified model often presents itself as a rational response to a challenge. It points to a set of practical and conceptual tools with which the typical Hedgehog feels at home. The outcome may descend into “group think”, in which orthodoxy comes to reign in a company, an industry and even a nation. It becomes not merely difficult to challenge the narrative, but actively dangerous for careers to do so. History presents us with many examples where people lost a great deal more than their job for questioning orthodoxy.

Hedgehogs tend to dominate companies which have poor narratives. They deliver results, but do not raise disturbing questions. Indeed, they abolish such questions by shrivelling the model and the values to a “nothing but”, and reduce the narrative to a repetition of a few catch phrases. Foxes cannot thrive in such an environment, and the essential contribution that they make to renewal is excluded. There are, of course, many other sources of error to which we can point. People armed with clear models and good data, clear values and a well-propagated narrative still make dreadful errors of judgement. Some of these are plainly “happenstance”, whilst others are due to two common errors in model building. It may be helpful to explore these.

We tend to believe what other people assert to be true. There have been large numbers of studies in which people changed their minds when exposed not to new data, nor to different ways of thinking, but merely to majority views. This remains true even when the majority is plainly wrong: when the horse is black and they all swear that it is today. We also tend to believe what we want to believe. That is, if we believe that the Internet is immune to economic logic, that ours is a special destiny to conquer and civilise, that the South Sea Company has the ability to generate money from air and rhetoric, then we blow ourselves a bubble. If everyone else seems to believe it as well, then we become ever-more whole hearted in the pursuit of soap-film delusion.

Markets go through crazes of this sort: Japan, emergent markets, junk bonds, service companies, technology companies, the Internet, China, derivatives, hedge funds. Markets go through what might be called ‘prescription’ crazes in much the same way: essentially, the belief that there is a managerial magic bullet and that those companies which have fired it are especially blessed and insightful. No company chair could appear before analysts in the early 1990s without mentioning re-engineering, down-sizing, out-sourcing, globalisation; and five years later without dotting their coms; or making reference to CSR (corporate and social responsibility) or China. As these things are usually grafted onto the organisation without any meaningful blood supply, they are usually at best distractions and at worst blocks to more structured debate.

Management teams which have lost their narrative are also susceptible to the madness of crowds. If the admired competitor has adopted a financial measure or an organisational strategy, then they will feel bound to follow. Consultancies and investment banks invest a great deal in creating and then marketing such measures. There are often confluences of interest, as with the re-engineering enthusiasm of the mid-1990s. Hardware vendors and software people, consultants and corporate IT specialists all found that this rationale allowed them to sell capital projects to senior management. Much the same was true of the Y2K pseudo-panic.

Models are created from an understanding of what matters to an organisation. Crude models connect “customers” and “costs” to “profit”. Subtle models understand what customers want, and find flexible and effective ways of surpassing their expectations. False models take these grafted on portmanteau words and link them up what does in fact matter, often masking areas which need attention. The force sent to pacify Afghanistan was seconded people to help Afghan tribesmen come to grips with gender sensitivity. The Russian invaders sent commissars to enforce communist orthodoxy amongst the troops. These often countermanded orders from military authority, and maintained a separate command chain to Moscow. Several value systems were plainly running in parallel for these choices to have been taken. A single synthesis had not been forced out of the system. That is, action preceded clarity. If we return to the McKinsey study about manufacturing discussed earlier, in which around half of all costs were incurred in “project definition”, then here we see exactly such costs expressing themselves in poorly defined projects.

Organisations abound with busy-work complication makers, to whom Hedgehogs are often a welcome antidote. However, in the absence of clarity – and when senior management emit statements about anything from emergent markets to sustainability without context or follow-through, such individuals have an open invitation to graft their preoccupations where they choose. State activities, in particular, tend to generate norms for each other to comply with. Each of these is no doubt worthy in its own terms, but collectively they can generate sclerosis.

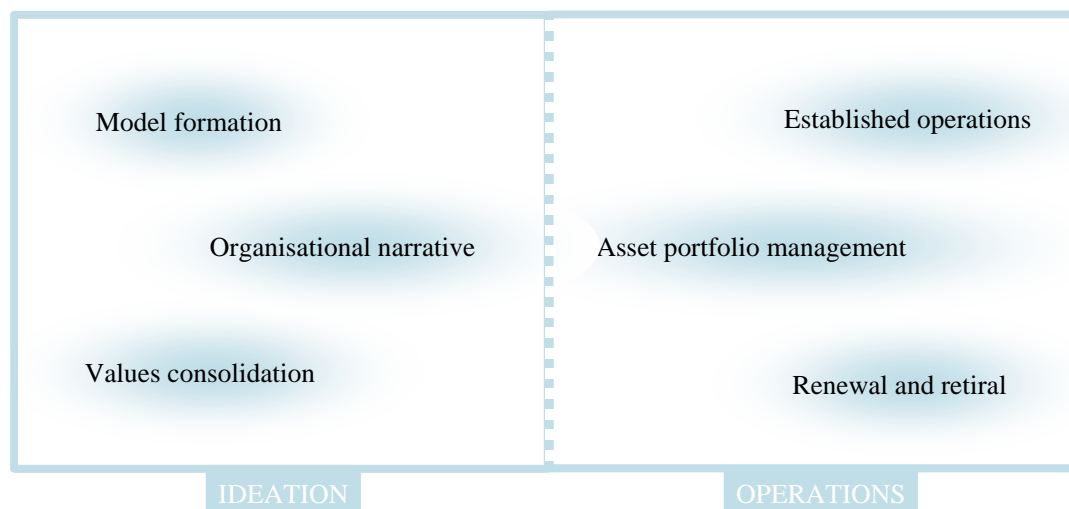
On how complicated organisations could improve their decision-taking

Figure 6 below shows two clusters of activity. On the left, labelled “ideation”, we find the three factors on which most of this paper has so far dwelled. That is, we find value definition, model creation and the development of the organisational narrative. All of this underlies what is conventionally termed “strategy”, a terms which has, unfortunately, come to encompass both everything and nothing. In brief, these are the processes by which an organisation comes, collectively and by iteration, to an understanding of how its operating environment works. It defines for itself what it does and does not want to have happen. It finds a way of articulating all of this in an accessible manner, so that day-to-day operations and option generation are steered by these insights. That is, they generate the model, the values and the narrative. The box on the right will be, for many, rather more familiar terrain. It contains the central focus of attention in most organisation namely, operations. It also contains the activities that lead to renewal and retreat; the processes by which new things are started, and through which established but now unwanted things are retired.

Portfolio management systems arbitrate over the flow of scarce resources. The future flows of funds and other resources are reviewed in the light of the projected needs of established and potential projects. Choices are made about their allocation. Should the organisation support this option or that one? How should projects be ranked? How does one compare dissimilar projects? Should the organisation seek external resources, use only what it can generate internally, or actually reduce its asset base in the form of asset sales, redundancy, dividends and so forth? The portfolio management system always arbitrates amongst aggressive operational rivals for scarce resource. Quiet voices are often silenced, and the major entities capture disproportionate resources. Many studies have shown that capital allocation in large companies is predicted almost

entirely by the amount of capital that a division already holds, for example, and not its profitability or its perceived long term prospects.

Figure 6: Ideation, strategy, operations, renewal and retiral;



As we have already seen, however, it is almost impossible to undertake either renewal or rational asset allocation without giving attention to the left hand box in Figure 6, labelled “ideation”. Organisations which have not undergone the ideation process properly, however, often develop a series of pathologies. We have discussed these in the preceding section. At least some of these – over-focus, rule by Hedgehog, diversion into crazes and enthusiasms, iteration of the past – mask the absence of clarity from themselves, if not from external scrutiny or from events.

Figure 6 shows that a barred line separates the two boxes, symbolising the somewhat vague connections which lie between these two sets of processes. These are uncertain for two reasons. One is that we simply do not understand how groups of people in fact do achieve clarity. Equally, however, what we do understand about this points to social, iterative, parallel, indirect procedures, such as those which characterise politics, small community life, kinship and the like.

Our everyday lives are more governed by tacit rules that we have acquired throughout our lives than they are by formal calculation and abstract legal structures. This is the ‘narrative’ that we have discussed earlier. Organisations also acquire collective insight, experience and common sense, also embedded in social interactions which are a group narrative. However, organisations are trying to do much more complicated things than are most small communities, and they are changing much faster. They do not have the cultural continuity that is enjoyed by the child, nor do they have the sheltered time that is granted by childhood in which to learn. This suggests two rather separate points that need to be taken into account when designing processes that are going to generate fresh insight:

- One, the features that we have discussed and placed in the “ideation” box have to be nurtured and actively developed if the organisation’s potential for clarity, agility and renewal are to be tapped.

- Two, the ideation box itself does not at all exist in isolation. It has to be connected into the information flows which drive the rest of the organisation.

The consequence for a process designer is that the two boxes are conceptually, but emphatically *not* operationally, separate. There are flows of knowledge that need to be created and directed, and activities which consolidate and in other ways make the information which is generated accessible to those who need it. They also create structures that predispose the organisation and the individuals in it to behave in certain ways. They help to orchestrate the formation of a narrative.

The remainder of this paper is concerned with the architecture and nature of these flows. As we have already mentioned, the detailed requirements of an individual organisation depend on its scale, life cycle and so forth. How a group of technical enthusiasts who are bound together in common cause in a non-governmental organisation should best conduct their dialogues is, of course, different from the best fit that can be created for a huge industrial organisation, embedded in the end game of its industry. That said, both groups need to undergo the same information exchanges, the same processes of rationalisation into a narrative, the same connection into operational matters, the same review of the outcome of all of this. Let us call these structures “ideation architectures”.

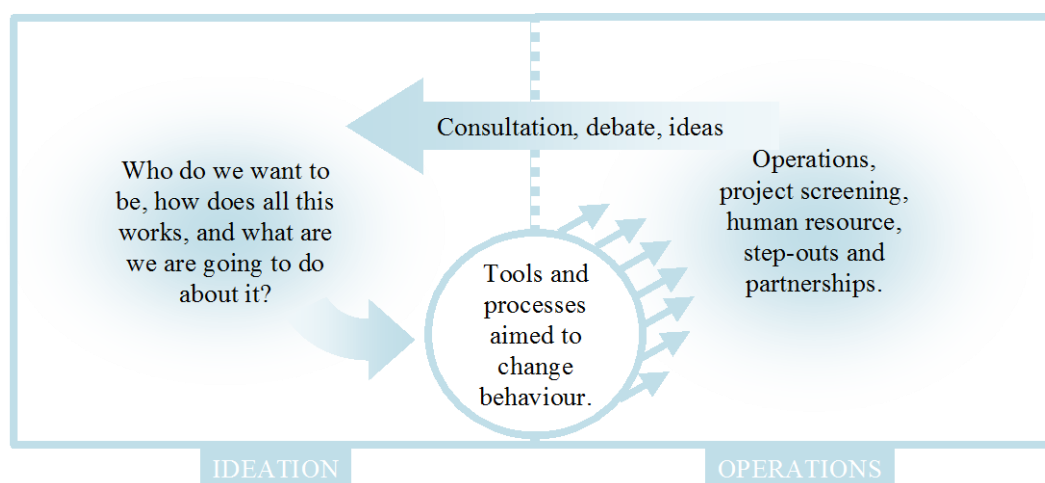
Ideation architectures

Interventions that create useful reactions need to be specific, both in the sense that they are designed to elicit a specific response and also that they are targeted on specific activities, operational units or individuals. At the same time, organisations need predictability, in the sense of there being regular cycles of reporting and analysis, proposals for action and decision mechanisms that evaluate these. Naturally, the “operations” box will contain elements of both of these, but whilst numerical analysis and reporting are normally extremely well-developed, the activities that connect all of this to the ideation side are usually weak, informal or non-existent. The generic connection between the two boxes is sketched in Figure 7.

While this all sounds rather abstract, it has major operational implications, particularly at board level. Boards have a variety of roles, such as representing the interests of the shareholders, exercising statutory controls and overseeing operations. They are also responsible for creating direction in the often loosely-coupled units of activity within the organisation.

We have already discussed the growing pressures under which organisations now work. Boards have been forced to respond to this intensification and have, on the whole, done so in much the same way. The drive has been to simplify the organisation, both so as to manage it in an ever-more complex environment, but also to be able to “explain” the organisation to stakeholders. Public sector organisations have followed similar steps, with particular emphasis on cost control and accountability. Board members have begun to act more as individual line managers and advocates for their area of responsibility, rather than as the unified team to which line-managers report.

Figure 7: Generic connection between ideation and operations



The requirements of Figure 7 question the adequacy of this response. Indeed, the weight of the analysis suggests that this pattern of response is likely to reduce adaptability and innovation, and increase the likelihood of herd behaviour, iteration of the past and the rule of the Hedgehog. There are, therefore, strong reasons why boards should give more, rather than less of their time to generating clarity around the model, values and narrative of an organisation; and identifying options for change. However, we have also seen that boards cannot take on the Ideation part of Figure 7, and leave operations to their line managers. The generation of sound insight requires the full range of insight and experience that the organisation possesses.

This casts the board into a new light: as the *enabler* as much as the director of events. Naturally, there are many examples of fine leaders whose individual skill lies in just this field. They induce people to speak together in ways which go beyond immediate issues. They encourage by listening those who have new ideas to bring forward. Their impact is clear and distinctive. However, boards which collectively act in this way are extremely rare; and those which set out to do this in a purposeful way, rarer still.

If the board needs to change its style and emphasis, it also needs a cadre of support staff who can design processes and run them, undertake analysis and present it effectively and, in general, fulfil the requirements of specificity with which this section began. Such people were once thought of as ‘head office staff’, perhaps as ‘planners’, and treated with great wariness by operational units. This, to, will have to change. It may be that people in the operational units should have a part of their time dedicated to such activities as a legitimate and normal part of their daily work load. Smaller organisations do this as a consequence of their size and, where something of this sort has been implemented, seem extremely effective at it.

How is all of this to be made operational? The answers lay in the process: in defining what modules of activity need to be done, by whom, in which order, with what output, and used in what manner. As we have already noted, processes work best when they operate under a predictable time-table. The output from one module then feeds into the starting position for its successors. Process owners thus have the responsibility to

their customer-consumers for their product, and the entire system can be encapsulated in appropriate systems of evaluation and reward.

Large organisations will have regular reporting cycles. They will have quarterly results that build to annual returns. They will have an annual process that gives unit managers targets and assets. The more sophisticated will have annual or biennial source-and-disposition reviews, by which cash and other resource flows are projected forward a few years. Some will have road maps for new projects that they intend to implement. Often, however, these sit in isolation to each other or to any attempt to integrate them with “strategy”, or with the understanding of the operating environment that we have called the “model”.

Plainly, a much better cyclical process asks its questions in broadly the right answer, and arrives at its answers in ways that support decision processes. For example, the development of the model of the operating environment needs to precede attempts to say what response the organisation could, should or will make to this. Models need to be developed with care, and it is not uncommon for this to be run on a two or three year cycle. Scenario planning is a good example of a process which undertakes this. It starts from the lessons learned from a previous cycle, for new concerns and areas of white paper, where the organisation feels itself ignorant. Many commercial organisations are concerned to understand “China” better, for example, in the sense of what it may mean for their interests in the medium term. Many do not even know what questions they ought to ask.

The scenario process tries to coax out the key variables, thus defining a space in which this particular aspect of the future may develop. For example, it might turn out that the socio-political issues which seem deeply important to the future of China itself are much less important to the organisation undertaking the study. China, it turns out, is a metaphor for the hundreds of millions of skilled, disciplined low wage workers who are becoming accessible across the world, and what they may do to issues such as quality control, market differentiation and local attempts to control imports into the place where the organisation wants to sell such products. The relevant space is spanned by dimensions about the intensity and openness to control of this process, and by socio-political responses to these events in the organisation’s key markets. The scenarios then track events through two or three plausible and important ways in which this dynamic might work itself out.

You may note a number of interesting things that are happening when this is done:

- First, the model is tested and changed until it becomes relevant, believable and useful. Vague generalities are turned into specific issues to which detailed attention can be given. Acrimony can be overcome, for contentious current issues are set in the context of impersonal forces and the grand sweep of events, trivialising individual disagreement and ambition.
- Second, a great deal of thought is condensed into an accessible form. Readers may recall the distinction between declarative and procedural knowledge: knowing how bicycles are ridden and having the skill to ride one. Stories that are embedded in a model act to transplant that insight into an audience’s head, where it resides as procedural knowledge. The suddenly know something about how to navigate these hitherto uncharted waters in a deep, intuitive manner.

- Third, organisational priorities can also be embedded in the analysis. For example, the organisation might have started the analysis by indicating concern about quality control management in diffuse and duplicated supply channels. This is then built into the story that is told as a key element on which the audiences mind should linger.

Model building is a relatively unthreatening pastime. Assessments of **value systems** are much more difficult to undertake from a cold start. Indeed, it can be difficult to engage people on this topic when things are going well. People will begin to discuss these issues when there has been a moral panic, an inexplicable public affairs or market-related crisis. In the sense that they are used in this paper, “values” refer to choices about operational style around which there are no clear objective figures of merit. An audit of the *de facto* choices exhibited by an organisation in its operations may show up unresolved questions, vague fudges and outright contradictions.

The simple fact of having assessed and characterised the values of the organisation is a powerful incentive to take this work further. There will be obvious gaps, contradictions and choices which point in the wrong direction. Characteristic questions arise which the organisation may never previously have debated, and new conversations will start that enrich the narrative. Typical questions for multinational corporations may include questions of the following kind:

- Should we have a global policy to labour relations?
- Is it even conceivable that we would not have a universal policy on health and safety?
- How are we to operate in environments in which some form of corruption is the norm: indeed, can we operate in them?

Any question of this sort – whether pertaining to the model or the value system – can be given relatively formal definition and assigned to a specific group of people to develop. On occasion, it can be outsourced to expert groups, and outsiders can, where appropriate, be invited to participate in such discussions.

The following process format can be extremely helpful. Following on from a general, perhaps diffuse, overview that has nevertheless been endorsed from the heights of the organisation, an informed group are asked to bring this down to between five to eight clearly defined questions. This is done by a simple process in which ideas are tabled, clustered and the resulting clump appropriately named. These clumps are then ranked for their importance and urgency. The items receiving top priority in this ranking are assigned to teams for further development. The issues arise from senior management’s endorsed view of what matters, have been selected by a reliable cross section of the organisation to be important and urgent, and are thus presented with considerable internal legitimacy and to equivalent interest. The entire activity is intended to define and explore the new knowledge that the organisation needs in order to take better decisions.

Formal process and general good will are likely to take an organisation only so far. These need to be reinforced with more pragmatic incentives. Large organisations do *not* respond well to broad brush directions; change is driven much more effectively by a myriad of small measures which taken together, create a “wind” that blows in the right direction.

If an organisation wants more innovation, for example, it is virtually useless for the board to demand this directly. Rather, analysis has to identify the individual pressure points which help or hinder the specific kind of innovation which is required. Steps are taken to adjust these in the desired direction. This might, for example, affect policies around recruitment, selection, development and reward. It might promote work styles which are known to enhance innovation. The status of individual innovators in the organisation might be raised in a wide range of ways. This approach opens up an extraordinary sweep of possible actions. For example, independent studies in European car manufacturers and pharmaceutical companies have shown that the physical layout of work spaces, or the way in which after-hours socialisation is or is not encouraged have considerable affects on the development of new ideas.

The most straightforward incentive on operational units to think broadly is to require that they embed their proposals within the broad issues which face the organisation. For example, if scenarios have been used, then the proposal must be robust in all of them. Issues such as environmental or social impacts must be explicit. The decommissioning process is factored in to the project accounts in ways that take account of prospective changing legislation. The proposal must respect the values to which the organisation has decided to adhere in more than a cosmetic manner. Project screening criteria – discount rates, the criteria used – shape the future organisation. Such methods are, therefore, extremely powerful in managing that shape. If they are not used, coherence will emerge only by accident.

Conclusions

Organisations make mistakes for two main reasons. First, they may have perfect and widely-shared clarity on the situation which faces them, and still be surprised. There are ways of mitigating such risk, but it is an essential part of the world in which we live and errors of this sort are *not* always avoidable. The second type of error can be avoided. It occurs when organisations lack clarity about how their operating environment works or when, for a variety of (correctable) reasons, they are unable to respond to this insight, choose to ignore it or to deny its validity.

The strength and multiplicity of the pressures on organisations will increase as the world becomes more competitive, more closely-coupled and more complex. The likelihood that they will fall into one or the other type of error will therefore also increase. If they follow fashion/trends, and go where their peers lead, they will sink into commoditisation. Differentiation that is based on insight and innovation is the *only* counter to this fate. We also note that the same source(s) of clarity that helps to avoid error is also a key element in the response to commoditisation. Organisational renewal consists, therefore, of harnessing the immense amount of knowledge and skill that is available within and outside the organisation so as to come up with a distinctive response that suits the new conditions. Such steps used to be episodic and in response to events. Plainly, however, organisations now need to make this a continual process, and to lead rather than than trail events. This needs to be organised, resourced, overseen and implemented with wisdom.

