



University of Pennsylvania
ScholarlyCommons

Management Papers

Wharton Faculty Research


12-2004

Crossing Epistemological Boundaries: Managerial and Entrepreneurial Approaches to Knowledge Management

Max Boisot

Ian MacMillan
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/mgmt_papers

 Part of the [Business Administration, Management, and Operations Commons](#), and the [Strategic Management Policy Commons](#)

Recommended Citation

Boisot, M., & MacMillan, I. (2004). Crossing Epistemological Boundaries: Managerial and Entrepreneurial Approaches to Knowledge Management. *Long Range Planning*, 37 (6), 505-524. <http://dx.doi.org/10.1016/j.lrp.2004.09.002>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/mgmt_papers/162
For more information, please contact repository@pobox.upenn.edu.

Crossing Epistemological Boundaries: Managerial and Entrepreneurial Approaches to Knowledge Management

Abstract

It is possible to identify two distinct yet complementary epistemological paths to knowledge development. The first one is holistic and field dependent, and builds on the concept of plausibility, and we associate this path with an entrepreneurial mindset. The second is object-oriented and builds on the concept of probability; this path can be associated with the managerial mindset. We believe that both managerial and knowledge management practices have emphasized the second path at the expense of the first. To restore the balance, knowledge management needs to develop processes and tools – associated with scenarios and real options – that will allow it to operate credibly in possible and plausible worlds, so as to extract value from them. We propose a systems framework for thinking through the nature of such tools.

Disciplines

Business Administration, Management, and Operations | Strategic Management Policy

**Crossing Epistemological Boundaries:
Managerial and Entrepreneurial
Approaches to Knowledge Management**

Max Boisot

Universitat Oberta de Catalunya

Ian MacMillan

The Wharton School, University of Pennsylvania.

July 28th, 2004.

To be published in *Long Range Planning*

Abstract

We identify two distinct yet complementary epistemological paths to knowledge development. The first one is holistic and field dependent and builds on the concept of plausibility; this path we associate with an entrepreneurial mindset. The second is object-oriented and builds on the concept of probability; this path we associate with the managerial mindset. We believe that both managerial and knowledge management practices have emphasized the second path at the expense of the first. To restore the balance, knowledge management needs to develop processes and tools that will allow it to credibly operate in possible and plausible worlds, so as to extract value from them. We propose a systems framework for thinking through the nature of such tools.

1.Introduction

We are interested in the evolution of knowledge management as a sub-discipline of management. Specifically we are concerned with how knowledge-bearing agents (individuals or groups of individuals with similar knowledge) develop and deploy to create value for themselves, or in other words manage their knowledge. Much of what today passes for knowledge management has its origins in practice, and *qua* practice, knowledge management has not much bothered with epistemological issues. Yet without secure epistemological foundations, knowledge management is unlikely to evolve from a practice into an intellectual discipline. In this paper we argue that knowledge management is open to multiple epistemologies, which give access to alternative types of knowledge worlds, each with characteristic challenges and opportunities for knowledge management: possible worlds, plausible worlds, probable worlds and actual worlds. Knowledge-bearing agents (hereinafter “agents”) use their knowledge in these worlds to take actions that secure future resources. The challenge for knowledge agents is to deploy scarce their resources in each of these knowledge worlds in such a way as to secure sustaining resources that exceed the resources committed. A key challenge for *knowledge management* is to understand the nature of , and the boundaries between, these worlds and, by implication, how and under what circumstances they can appropriately be crossed by knowledge agents. This paper explores the nature and boundaries of these worlds and shows how they can be construed as sources of opportunities and of value for those who know how to move within and across them.

Consider the following investment events:

- In the Internet bubble of the mid 1990s, small firms raised huge sums of venture capital in large numbers. Many came to market with little or no track record to speak

of and never having made a profit. Investors placed their money in these firms on the strength of what their rudimentary business plans promised and of a blind faith in the future of the Internet. The bubble burst in 2000 and the irrational exuberance of the market for these online companies, even ones with real potential, vanished with it.

- One day, before going on a trip to the United States, Masaru Ibuka (then Honorary Chairman of Sony) asked Norio Ohga (then Executive Deputy President) for a simple, playback-only stereo version of the 'Pressman', the small, monaural tape recorder that Sony had launched in 1977. He wanted to be able to take something light and portable with him on his travels. In 1979, Sony launched the 'Soundabout', a personal stereo that was later relabelled the Walkman. It was developed on the basis of nothing more than a strong personal hunch. Sony expected to sell 5000 Walkmans a month. Within two months of the product launch it was selling ten times that amount and has since become a cultural icon.
- Inside large organizations, innovative managers approach their board with meticulously drawn up business plans and a mixture of objective statistical facts and estimates to justify these plans. The managers proposing these plans know - and often, the members of the board know - that these highly uncertain estimates are tentative and thus very likely to change as events unfold. Nevertheless, based on the evidence provided by the figures the board will decide whether or not investment in these proposals is justified. Entrepreneurs persuade venture capitalists to invest with perhaps even more uncertain business plans.
- To be sure that they can meet their obligations, pension funds are required to place a significant proportion of their resources in risk-free investments. For this reason they

invest part of their cash in money market instruments of proven reliability. This is a highly liquid form of investment that yields low returns but is certain to give back the cash originally invested.

In pursuit of an appropriate knowledge management perspective, in what follows we conceive of *knowledge as comprising a set of beliefs which informs decisions by agents to take actions that consume the agent's (scarce) resources*. With this conception of knowledge, each of the above cases involves deploying knowledge: that is, it involves taking action that consumes the resources of agents on the basis of sets of beliefs that these hold individually or collectively. In the first example, the beliefs were vague and carried little or no justification and the risk was significant. In the last example, the justification is well established and the degree of certainty is high. In the second example, what mattered was the strength with which the belief was personally held by a powerful decision-maker, irrespective of whether it could be justified to outsiders. And in the third example the key requirement was to justify the belief to members of the board, whether or not one held it oneself. In each of the above examples, action is a commitment of resources to belief in a perceived opportunity. In the four cases, beliefs are more or less strongly held, and more or less capable of being justified to outsiders. The opportunity is characterized by high or low levels of uncertainty that gets eliminated with the passage of time. In a world characterised by complexity, variety and uncertainty, the challenge is to take actions that are appropriate to the nature of the belief that is held, and to match the latter to the level of uncertainty that inheres in a given situation

Such a matching exercise illustrates Ross Ashby's Law of Requisite Variety. This states that the variety of stimuli impinging upon a system must be countered by the variety of responses that the system can muster. As Ashby put it, *only variety can destroy variety*.¹ Yet

some variety constitutes noise for the system and therefore calls for no response by the system save that of filtering it out. A system that is incapable of filtering out noise from the set of stimuli that it responds to is condemned to dissipating its scarce resources unproductively as it overreacts to every opportunity or threat, real or imagined. Variety reduction then becomes necessary in order to filter out stimuli that do not give rise to actionable beliefs – i.e. to a form of knowledge. However, in intelligent systems, that is, systems capable of forging meaningful representations of the states of the world that they respond to, what constitutes noise and what constitutes information for them will itself be a function of their models of the world – i.e., the prior beliefs that they apply to the interpretation of stimuli.² And the larger the number of models that the system can choose from, the wider the range of possible interpretations of stimuli available to it. Also, by implication, the larger the number of models that the system can choose from, the more tentative and uncertain become the filtering processes that distinguish information from noise.

The law of requisite variety is a call to *action*, and knowledge is an essential ingredient of effective action.³ But what is likely to constitute *sufficient* knowledge to take action? How do the different types of belief that an agent is willing to act upon relate to each other? And how do they increase in certainty? Furthermore, intelligent action is action that can handle variety adaptively within a given time frame – *requisite* variety has a time dimension. How, then, to zero in on the relevant models and appropriate beliefs in a timely fashion?

Getting answers to questions such as these is likely to grow in importance in the coming years. The questions can be subsumed under a broader one: how might we *economically* manage our scarce knowledge resources under conditions of uncertainty? The burgeoning field of knowledge management has hardly ever framed its challenges - let alone

attempted to answer them - in these terms. Why so? We argue that to be able to answer such questions, knowledge management will need to develop its epistemology – for our purpose, the question of what constitutes valid knowledge under different circumstances. Epistemology provides the basis for action and thus serves as a foundation for the institutionalization of practice. Absent a credible epistemology, knowledge management is playing Hamlet without the Prince, condemned to remain a loose collection of empirical practices rather than evolving into a full-fledged intellectual discipline.

Epistemology, however, is not physics. Its principles and its application will vary according to time and place. There is therefore a need to identify the different circumstances – social, technological, economic, etc - in which knowledge is considered valid and actionable. In this paper, we unpack the Platonic view of knowledge as “justified true belief” into combinations of its different components. We argue that different mindsets will emphasize different combinations of the components “belief”, “truth” and “justification”.

We shall proceed as follows. In the next section 2, we introduce some of the epistemological issues that confront knowledge management and put forward a simple scheme for relating these to each other. In section 3 we provide a brief overview of the current state of knowledge management as well as of its antecedents in the institutions of science. In section 4, we show the relevance of our interpretive scheme for the current challenges that confront knowledge management. A conclusion follows in section 5.

2. Defining Epistemological Boundaries

Epistemology – further discussed in our appendix - is the study of the nature of knowledge and justification.⁴ Plato, aiming at the attainment of certainty, an infallible state of mind,⁵ took knowledge and uncertainty to be antithetical to each other. In the *Meno* and the

Theaetetus (c. 400 b.c.) he defined knowledge to be *justified true belief*. This definition identifies three individually necessary and jointly sufficient components of what counts as infallible propositional knowledge: a truth condition, a justification condition, and a belief condition.

While it is appropriate for philosophical debate, when it comes to resource allocation decisions the Platonic definition of knowledge, which demands true justified belief, is too restrictive. More crass than Plato, we are interested in knowledge that informs agent **action**. We suggest that it is beliefs that underpin agent action – so an agent can have justified true beliefs; justified beliefs; true beliefs; and, unconstrained beliefs. Thus different kinds of knowledge are possible, not all of which can be expressed propositionally and not all of which require the presence of all the three components for action-based knowledge.

The naturalistic perspective on knowledge, for example, concerns itself with possibilities for action in which knowledge, as well as being representational, can be tacit⁶ or embodied in skills and know-how⁷ – what the Greeks termed *techne*. From this perspective, one drops the truth condition and settles for justified belief as a ground for action – if it is tacit or embodied, such belief may not take an explicitly propositional form and nor will its justification. If, on the other hand, one drops the justification condition, an agent may settle for true belief alone – i.e. an agent may be willing to act on beliefs that square with its own prior experience without seeking to justify such beliefs to others. Finally, an agent may settle for belief *tout court* and act on a ‘hunch’ where neither truth or justification are involved. In the last two cases, however, a price is paid in the form of a loss of social legitimation of the belief in question. Those agents who ‘have the courage of their beliefs’, however, are often willing to pay the price.

As we have said above, the challenge for agents is to deploy scarce resources in each of the knowledge worlds in such a way that the agent secures sustaining resources that exceed the resources committed. Such sustaining resources are only secured in actual worlds – ie, those characterized by certainty. In worlds characterized by varying degrees of uncertainty, however, there is some advantage in being able to develop appropriate forms of knowledge and to be able to migrate from one world to another in line with changing levels of uncertainty.

Thus, we can take the three components of justified true beliefs as representing different types of constraints on our definition of knowledge and then go on to rank different epistemologies by the degree of constraint that they impose upon us. Clearly, the most demanding constraints on our beliefs are that they should be both justified and true. As we go about our daily business, few of our beliefs are ever actually called upon to pass this demanding test and, indeed, few, if any, of our daily actions are based upon such epistemologically demanding conditions. We can greatly expand what we will count as valid knowledge, therefore, by dropping the truth condition and settling for a justification condition on its own. That is the strategy adopted by science. It is constrained by the requirements of justification – an essentially *social* process based on the authority of sensory evidence and logic rather than on charismatic or institutional authority - but no more than this.

Alternatively, we can drop the justification condition, acting solely on the basis of true belief – a strategy that gives more scope to personal intuition, conviction and experience. Of course, this type of knowledge often being highly subjective⁶, we may be unable to get others to accept it and we may then end up acting alone. Finally, we can have a large number of beliefs that are unconstrained by either a truth or a justification requirement but that we are

still willing to act upon, albeit in a more cautious and tentative fashion – i.e. we hedge our bets.

We represent our constraint-relaxation model in the form of a Venn diagram (see Figure 1). The diagram has two intersecting circles set inside a larger circle. The larger circle, 1, represents the full range of what we can take to be knowledge – i.e. beliefs that to some degree we are willing to act upon. The two intersecting circles, 2 and 3, represent the more restricted views of knowledge, i.e. those that either meet the truth condition, the justification condition, or both. Clearly, the most constraining condition is to be found in the area given by the intersection of circles 1 and 2. The least constraining one is given by belief on its own – i.e. Circle 1 minus circles 2 and 3. We briefly discuss each condition in turn:

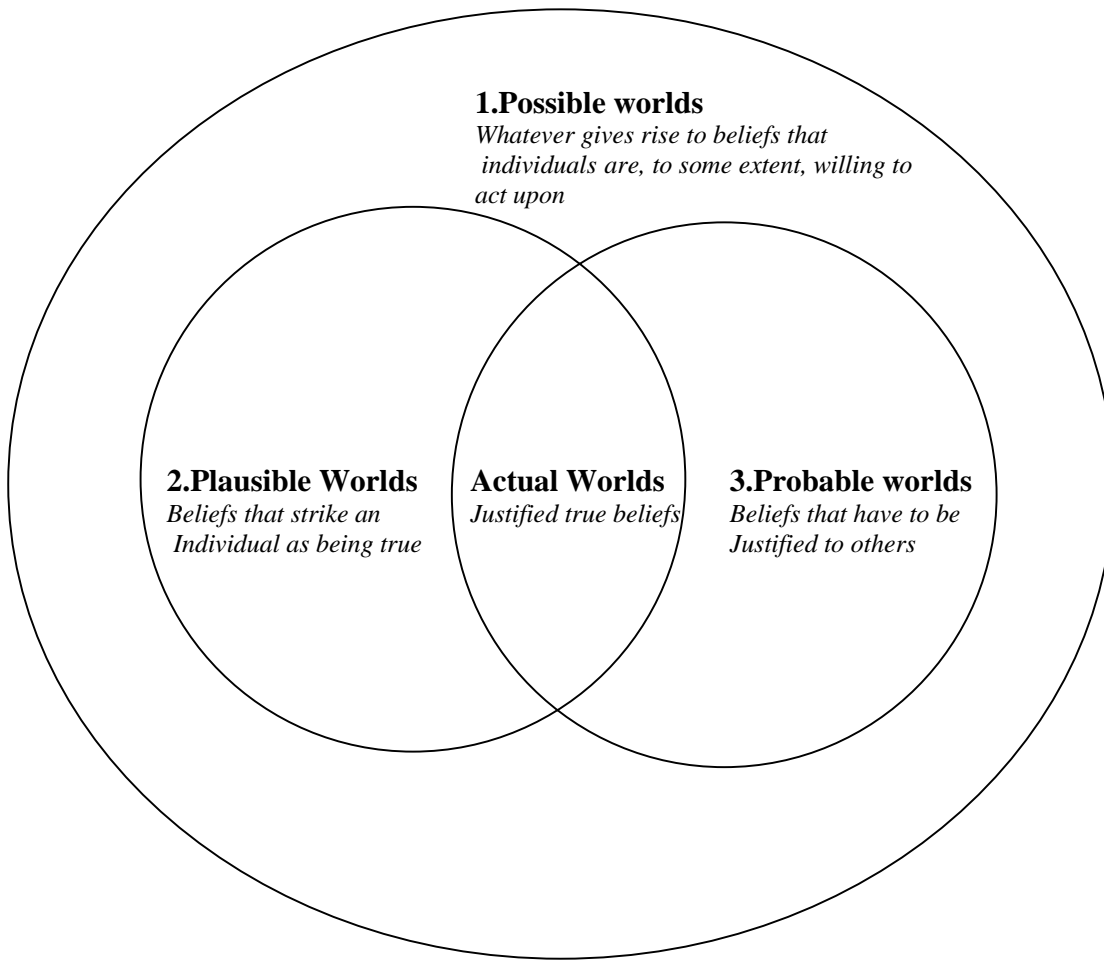


Figure1: Possible, Plausible, Probable and Actual Worlds

Circle 1 minus circles 2 and 3 describes beliefs that are constrained neither by truth nor justification requirements. Here, ‘anything goes’⁸ providing that it does not contradict the laws of logic or of physics – i.e. providing that it is not actually impossible. Such unconstrained beliefs give us *possible worlds*. These worlds are characterized by phenomena to which no probabilities can (yet) be attached. This has two origins. The first is that a given

phenomenon may not give rise to discernible events that are amenable to measurement – i.e. the phenomenon will not be *identifiable* and hence it will not be possible to distinguish an event from a non-event. The second is that even if an event is so discernible, it may not be repeatable within an empirically relevant time frame and hence cannot give rise to a probability distribution.⁹ These have reproducibility as a core requirement. In short, possible worlds lack both the clarity and the regularity that underpin the formation of rational expectations. Here, neither memory, nor perhaps even perception, has much purchase. At the height of the Dotcom boom, a good number of start-up firms were sustained by little more than the possible worlds that they had identified for over-gullible or overly optimistic and greedy investors.

- Circle 2 describes beliefs that are constrained by a truth requirement but not by a justification requirement. Like justification, truth requires both coherence and correspondence with the facts. Yet in the absence of an external reference group, whatever strikes one as coherent and in correspondence with the facts – i.e. as being true - will remain personal and subjective rather than inter-subjectively validated. It may be based on deep intuitions and extensive personal experience but, unless an agent possesses a strong charismatic authority, it may prove hard to convince others of a truth without providing sharable evidence that would be acceptable to them. We are here in a *plausible world*, one which individuals might reasonably act upon but which lacks the objective (or inter-subjective) justification for collective action. It was inside such a subjectively derived plausible world that Masaru Ibuka of Sony initiated the design of what was to become the Sony Walkman. The decision could be justified on

the basis neither of reasoned arguments nor of observed regularities. Ibuka was able to act on the strength of entrepreneurial hunch alone.

- Circle 3 removes the constraint imposed by a requirement for truth and accommodates beliefs that have been justified. Where justification is not based on revelation or authority – we will not deal with these here - but rather on objectively demonstrable coherence and correspondence with the facts, such beliefs access *probable worlds*. These are worlds in which rational expectations can be developed and shared on the basis of the discernible regularities that reside in past experiences. They are accessible to probabilistic and statistical analysis, and allow action to be based on calculable risk. For this reason probable worlds can be justified to others. In our earlier examples, managers and entrepreneurs who lacked the resources to back their own hunches in the way that Ibuka did at Sony, had to justify their beliefs to board members or venture capitalists by presenting carefully analysed technical and financial data as evidence. We noted that they themselves did not have to believe in the truth of this evidence in order to put it forward as justification.
- The intersection of circles 1 and 2 accommodates Plato's definition of knowledge, i.e. justified true belief. Such beliefs, justified to others either directly by the evidence provided by the senses or by the inferences that these allow, access *actual worlds* in which experience has immediacy and is indubitable. I do not doubt the existence of the laptop computer on which I am typing this paper; it is here, before my very eyes. Nor, typically, do I doubt the veracity of my bank statement, when it credits my account with a cheque that I paid in yesterday. The truth of my own experience squares with the justification provided by the bank statement to reinforce my belief that my account

has just been topped up by an amount corresponding to the value of the cheque. It is this kind of certainty that pension funds pursue in their investment strategies.

Each of the regions in the Venn diagram offers a distinct epistemological perspective that, to paraphrase Abraham Lincoln, will have validity to all of the players some of the time and to some of the players all of the time. One geometrical consequence of our Venn diagram is that both the number of plausible and of probable worlds is larger than the number of actual worlds, and that the number of possible worlds, in turn, exceeds the number of plausible and probable worlds. Actual worlds can be represented as singular outcomes that can be inter-subjectively agreed upon; plausible worlds contain singular outcomes that will not necessarily command inter-subjective agreement; probable worlds can be represented as a probability distribution of outcomes defined over a given range; possible worlds can be represented as a range of outcomes over which no probability distribution can be specified.

Each region in the diagram yields some kind of actionable knowledge even though only the intersection of circles 2 and 3 inside circle 1 yields the kind of knowledge that would have been acceptable to Plato or Kant. The knowledge in each region, therefore, has some value for action, but typically, any action in the real world will to some extent draw its epistemic resources from all regions simultaneously.

Furthermore, action, by generating new knowledge, will shift epistemic resources from one circle to another in both the inward and the outward direction. How, specifically, might we establish the value-for-action of these different kinds of knowledge?

In Figure 2 we take action as requiring some kind of a resource commitment that varies according to the world that an agent finds itself in.

Taking each world in turn:

Possible Worlds are characterized by unconstrained beliefs. An indiscriminating commitment to all possibilities would quickly deplete the agent's resources. Resource commitments to action therefore must be highly tentative and can only be based on the *option value* of a selected set of possibilities. Framed in terms of action, an agent might – just - get what it pays for. Scenario thinking in organizations exemplifies the kind of thinking required in possible worlds. Over time, some subsets of possible scenarios may acquire plausibility on account of their coherence; other subsets, with accumulating empirical evidence, may come to seem probable. In this way, from possible worlds an agent can gradually move either toward plausible worlds, toward probable worlds, or towards both simultaneously. The knowledge that inheres in possible worlds is generative of value, but it remains provisional. In possible worlds anything is possible – knowledge in this space comprises beliefs in possibilities - hints and hunches about what linkages among beliefs might be. However there is not yet any coherence of linkage among these hints, nor with the insight or experience of the agent. Furthermore, there is no history or evidence by which to corroborate and justify these hunches. The challenge for the agent is to recognize the potential for forging value from some subset of these possibilities by taking out options to exploit the value of whatever hunches can be navigated – either through plausible worlds or probable worlds (where profit potentials are created) into actual worlds, where value is captured.

Plausible Worlds: Knowledge here also has value, but now only for the agent (individuals or groups with common belief) holding the belief. In plausible worlds knowledge is true for the agent (but not necessarily for other agents) in the sense that the linkages among beliefs held by the agent have coherence with one another and with the agent's experience, in other words the set of beliefs is plausible to that agent. The challenge for the agent is to

recognize the potential for creating value from some subset of plausibilities by making “speculative” investments in these opportunities unseen by others. The agent thereby creates the potential to capture abnormal profits in the form of Knightian rents derived from the agent’s idiosyncratic insights. Expectations are sufficiently coherent and grounded in the agent’s experience to justify applying an agent-specific discount rate. The *subjective net present value* that results will reflect the application of the agent’s idiosyncratic insight - rather than objective criteria. In a plausible world an agent may strongly believe that it will get what it pays for, but, in the absence either of objective evidence to justify such a belief, or of a strong personal charisma, the agent will not easily persuade other agents to pay along with it.

Probable Worlds: In probable space knowledge is justified in that the linkages among the set of beliefs can be corroborated by external evidence. Through empirical testing and replication the outcomes can be replicated and a probability distribution assigned, thus creating socially justifiable probabilities. The challenge for the agent is to recognize the potential for creating value from some subset of probabilities by making risk-adjusted investments which create the potential for normal profits. The kind of replicable empirical knowledge available in probable worlds allows for the application of a socially derived discount rate; it thus has an *objective net present value*. Framed in terms of action, an agent will probably get what it pays for. Much scientific knowledge is of this type, not indubitable, but, on account of systematically recorded repetitions and replicated tests, highly *corroborated* and hence, highly probable.

Actual Worlds: In actual worlds knowledge comprises true, justified belief – either the agent is in a spot market where cash flows immediately out of spot contract or the emergent

successes of agents' investments in options, speculations or projects are cashed out. We argue, as William James did, that the certain knowledge available in such worlds to any agent which has direct or indirect access to the sensory evidence, has *cash value*.¹⁰ Framed in terms of action, the agent will certainly get what it pays for, and so will those who pay with it.

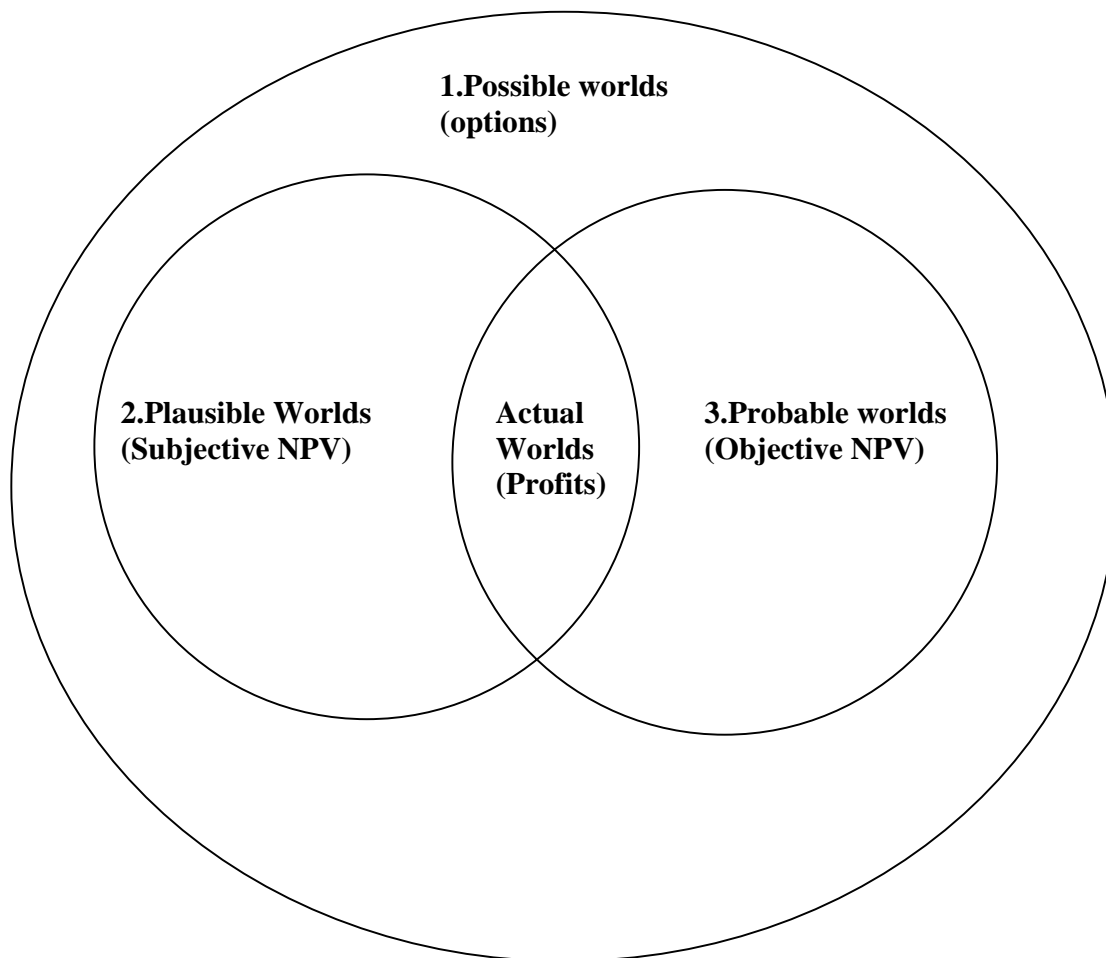


Figure 2: The Value of Different Worlds

The diagram highlights the fact that an agent has two quite distinct paths for navigating from possible to actual worlds:

- Via plausible worlds – an agent starts off building coherence into its beliefs thus moving from the possible to the plausible, and then seeks to establish a correspondence between these coherent beliefs and the real world in order to justify them.

- Via probable worlds – an agent first looks for justificatory phenomena in the real world that might correspond to its beliefs and then gradually builds up a coherent and plausible interpretation of this correspondence.

Plausibility and probability constitute alternative yet complementary bases for action.

Whereas the first bases action on the *coherence* of an experience – does it make sense? – the second bases action on the robustness of the experience’s replication – i.e. on its reliable *correspondence* to some recurring state of the world. Both help to underpin action, but each may kick in at a different moment on the journey from belief to action. In the early stages of any innovative process, for example, replication is hard to come by, if not impossible. For this reason entrepreneurial mindsets are more likely to look for plausibility and coherence before acting than managerial mindsets. The managerial mindset will look mainly to probability and correspondence to justify its actions. As experience accumulates over time and repeats, plausibility gradually gives way to probability as a basis for action.

If we think of a “mindset” as the embedded epistemology that an agent employs to navigate a path from possible worlds to actual worlds, there are two basic types of navigating mindsets. We will call the navigating strategy using the path through plausible worlds the entrepreneurial mindset. An entrepreneurial epistemology seeks to extract real world value by enacting plausibilities. On the other hand, the navigating strategy using the path through probable worlds employs what we will call a managerial mindset. An embedded managerial epistemology seeks to extract real world value by enacting probabilities.

Entrepreneurial and managerial mindsets will each draw on their respective epistemologies as a basis for action. Whereas the managerial mindset seeks evidence and

uses probabilities to justify its actions to others – stakeholders, shareholders, etc - typically, the entrepreneurial mindset need not. By contrast, whereas the entrepreneurial mindset has to have the courage of its conviction in its beliefs – after all, it typically has to “put its money where its mouth is” – the managerial mindset can afford to be more tentative in its beliefs before acting on them as long as it can justify them.

Now, as the uncertainties associated with the process of innovation are gradually removed or converted into measurable risk, and as the number of stakeholders associated with the process increases, the entrepreneurial mindset is required to cross the boundary between entrepreneurial and managerial epistemologies. Likewise, as firms experience the need to change and renew themselves – and hence to behave more entrepreneurially – the managerial mindset is required to cross the boundary in the other direction. There are costs and benefits, both personal and institutional, associated with such boundary crossing. Not everyone can make it.

These alternative mindsets and their associated navigation paths have profound implications for knowledge management. After a brief review of the current state of knowledge management we will look at these implications.

3. Knowledge Management: Between the Possible, the Plausible, the Probable and the Actual

As is often the case with emerging fields of professionalization, much of what today passes for knowledge management has its origins in practice – and in particular, in the spread of information and communication technologies.¹¹ And, as elsewhere, practitioners of knowledge management have not until now been much troubled by epistemological or foundational issues. This uncertain progression from a casual and empirical stance to a more

theoretically informed one is a well-established phenomenon.¹² Steel making in the nineteenth century, for example, was largely a matter of empirical trial and error that preceded the development of metallurgy as a science.¹³ Steel making, then, was understood empirically long before it was understood theoretically. But the kind of concrete knowledge embodied in such empirically derived practices was highly local and hard to replicate in a controlled manner. It lacked any capacity for *leveraging*, for getting extensively applied beyond the local context. The potential for leveraging is the great advantage that abstract science-based knowledge enjoys over more empirical practices, rooted as they often are in the vagaries of craft traditions.¹⁴ Such knowledge can have relevance and can be applied over a much wider and more diverse range of circumstances. The gradual application of science-based knowledge to steel making helps to explain why the total world output for steel grew from 500 000 tons in 1870 to a total of 28 million tons by the turn of the century.¹⁵

However, to date much of the knowledge that has been of interest to knowledge management has tended to be concrete and local in nature. It consists of rules of thumb, anecdotes, and best practices assembled and deployed within one organization and/or its related network, but not beyond it. Knowledge management thereby contributes to making better use of the knowledge that an organization already possesses. “Knowing what it knows” spares an organization the expense of ‘reinventing the wheel’. There is typically no intention of leveraging such knowledge beyond the confines of a single organization or network of linked organizations. The resulting knowledge “makes sense” - it gives them the right gut feeling to those who directly share the relevant experiences. But because the knowledge is local, replication and dissemination is a problem. Such knowledge is *plausible*, but only to those who directly experience it or to those who can take such experiencing on trust.

However, if knowledge is to move from small local populations to larger and more distant ones, access to the direct experiences that gave rise to such knowledge becomes increasingly problematic. How is another agent to trust such knowledge? How is a claim of having knowledge to be *justified*? To be justified, claims of having knowledge require independent and replicable testing. The results of such testing create knowledge patterns that follow distributions, and thereby allow resulting claims of knowledge to rest on *probabilities*.

Contemporary knowledge management to date has contributed even less to the *creation* of significant new knowledge. We say “contemporary” because there is one form of knowledge management that has been outstandingly successful in this respect. We refer to the practices of the scientific community.¹⁶ This kind of knowledge management has been around for some time. But the creation of knowledge by the scientific community does not follow the “logic” of knowledge creation as practiced in a corporate environment or, indeed, that advocated by knowledge management practitioners. The different ‘logics’ show up as differences in emphasis. In all three cases the key concern has been with control of the diffusion of new knowledge. Yet whereas the scientific community’s primary concern has been with the epistemological *validity* of the newly created knowledge, that of the corporate community has been with the economic *utility* of such knowledge.

The philosophy of science has a sub-branch – methodology - that takes the validity of knowledge as its central concern.¹⁷ It attempts to endow the practice of science with solid epistemological foundations. The more solid those foundations, the greater is the potential for leveraging scientific knowledge, both existing and new, and reliably extending its application to new fields. Arguably, it is precisely the absence of solid epistemological foundations that is undermining the numerous attempts at leveraging corporate knowledge creation both within

and across organizations. Yet because its primary concern has been with the utility rather than with the validity of knowledge beyond the boundaries of the firm, knowledge management as practiced by corporations has not so far felt any pressing need to secure its epistemological foundations. One plausible explanation for this, perhaps, is the daunting nature of the task: a two-and-a-half-thousand-year old debate on the nature of knowledge that goes back to Plato has not so far made those foundations any more secure. Drawing a boundary between useful and useless knowledge, therefore, may be easier for a corporation to do than drawing a boundary between valid and invalid knowledge.

So far we have argued that entrepreneurial and managerial mindsets arise in different circumstances and tend to draw on different epistemologies. As circumstances change, however, these mindsets not only need to cross epistemological boundaries, but they also need to be aware that they have crossed them. In the next section we briefly explore the nature of these boundaries and to show how they can be construed as sources of opportunities for those who know when and how to move across them.

4. Epistemology and Knowledge Management

From an epistemological perspective, the key skill in knowledge management involves understanding the basis on which an agent can move the products of its thoughts across the epistemological boundaries represented by the circles of our Venn diagram – i.e. from one type of world to the other. Both the worlds of the plausible and of the probable, for example, are sources of potentially fruitful hypotheses. If carefully analyzed and reflected upon, such hypotheses gradually gain in plausibility. On the other hand, if repeatedly corroborated through testing, they become more probable. In both cases they improve their epistemic status. Yet the world of the possible is large relative to the worlds either of the probable or of

the plausible, and when it comes to hypothesis selection, it offers an excessive amount of choice. An agent needs a good reason for selecting one hypothesis rather than another for further development and testing, since such activities consume scarce resources – after all, even options have a carrying cost.

In the world of the possible, if no prior distribution is available, no expectation can be built on the basis of recurrence. It must therefore be built on the coherence of the expectation relative to our other beliefs that act to constrain it.¹⁸ This is a theory-driven, sense-making activity that gives plausibility to an expectation and reduces subjective feelings of uncertainty. If, in contrast, phenomena exhibit recurrence and prior distributions are therefore available, an agent can subordinate the need for making sense of things to establishing a correspondence between prior beliefs and recurrent phenomena. Thus it will be the states of the world that suggest the appropriate trajectory to follow in moving from a possible to an actual world.

Given the lack of replicable precedents under conditions of innovation, the entrepreneurial mindset will tend to favour an initial move from possible to plausible worlds. It will then *enact* for itself the replications that will lead them into the region of Figure 1 where the plausible intersects the probable. It will do so by constantly testing prior assumptions against accumulating empirical evidence in a process known as *discovery-driven planning*. The managerial mindset, by contrast, focused as it is on the need to justify its epistemic stance, is generally more disposed to move into actual worlds via probable ones, and to build its theories on the basis of available empirical evidence rather than seeking out evidence in support of *a priori* theories that have been subjectively derived.

Society has a large say in establishing what gets placed within each of our worlds and what gets excluded, what constitutes legitimate moves across epistemological boundaries and

what does not,¹⁹ and what gets emphasized or played down. French culture, for example is more likely to stress coherence than Anglo-Saxon culture, which is more drawn to an empirical approach, and hence to correspondence. Chinese culture on the other hand will stress the complexity and multi-causal nature of phenomena in contrast to western cultures operating on the principle of Occam's razor and the quest for simple explanations.²⁰ Chinese culture is thus likely to draw the boundaries more loosely around our worlds than cultures reared in the Western scientific tradition.

Firms will often take their cues from established social practice in drawing boundaries. And for them as for other institutions, the issue will be one of balance. Draw the boundaries too tightly and you stifle innovative moves; draw them too loosely and scarce resources are squandered. The issue of balance leads us back to the question raised by Ashby's law of requisite variety. How much variety is actually requisite? Does every instance of variety call for a response? We can illustrate the nature of the issue by means of a diagram that presents Ashby's Law in a graphic form (see figure 2). The vertical axis of the diagram measures the variety of the stimuli to which a system is subjected. The horizontal axis measures the variety of the responses available to the system. Ashby's law locates adaptive responses on or below the diagonal in the diagram— i.e. the variety of a response at least matches the variety of the stimulus that provoked it. Yet, as we saw earlier, in a regime of high-variety stimuli, the sheer variety of responses that appears to be required might well lead to the disintegration of the system. At the other extreme, a system with little or no variety in its responses eventually fossilizes or gets selected out. (See figure 3)

The challenge for any living system, then, is to navigate between the twin threats of disintegration and unresponsiveness. Living systems endowed with cognitive capacities, however, have successfully evolved responses to *representations* triggered by the stimuli rather than to the stimuli themselves, that is, they draw on prior knowledge of the stimuli to filter out those elements of stimulus variety that constitute noise, concentrating their response on the much smaller variety of information-bearing stimuli that remain. In Figure 3 this more ‘cognitive’ strategy is indicated by the line AB’. In contrast with the horizontal line AA’, it does not attempt to match the variety of a given set of stimuli on a one-to-one basis with a given set of responses. Rather, through a filtering and interpretive process, it reduces the variety of the response called for by reducing the number of stimuli that it actually *needs* to respond to. The epistemological challenge for knowledge management may well be to help us decide how far down the scale we can legitimately move and still be responding to valid knowledge.

*The **behavioural** route responds directly to stimuli without filtering out the noise*

*The **cognitive** route filters out noise from stimuli and focuses on their representation*




Figure 3: Ashby's Law of Requisite Variety

Applying this thinking now to Figure 1, it is clear that moving towards the rim of the outer circle, 1, is a variety-seeking strategy that moves us up the vertical scale of Figure 2, whereas moving towards the intersection of circles 2 and 3 in the figure is a variety-reducing one that moves us down the vertical scale.

Looking at it another way, moving inward in the diagram from circle 1 to the intersection of circles 2 and 3 involves exercising options that already exist. Moving outward from the intersection to Circle 1, by contrast, involves creating new options. The movement inward, towards the intersection of circles 2 and 3, reduces uncertainty and renders knowledge more reliable and usable. Since it represents a movement towards increasing constraint, it will be incremental in nature. It can be associated with a process of *exploitation*, one that builds up and consolidates existing knowledge. The movement outward, by contrast, broadens the horizons of awareness and can sometimes lead to a radical restructuring of what has already accumulated in the inner regions. It can be associated with a process of *exploration*, the creation of new knowledge.²¹ Exploitation, then, is likely to stabilize and consolidate the existing knowledge base whereas exploration is likely to destabilize it and make it contingent. In Kuhnian terms we might say that inward movement entails puzzle solving whereas outward movement will often look more like paradigm creation.

Our analysis suggests that an effective knowledge management process would not only attempt to strike a balance respectively between outward and inward movements in Figure 1 – i.e. between variety-seeking exploration and variety-reducing exploitation - but it would also establish which of the different regions of the Venn diagram it was operating in at any given moment so as to match the behavioural strategy required to the epistemic resources available– i.e. it would align the relevant mindset – in our case, entrepreneurial or managerial – to the nature of the task and the knowledge available.

The foregoing has implications for the way we think about the process of innovation. Schumpeter placed innovation at the heart of the process of economic development. He viewed it not as the fruit of rational planning, but rather as the unleashing of ‘gales of creative

destruction'.²² Innovation, then, offers benefits to society, but it comes with a cost to some individual firms. A parallel problem exists for firms: new products and processes can disruptively supplant the firm's established order. Innovation thus produces growth, but it also produces losers as well as winners. Many agents inside firms would prefer not to risk being losers. They thus tend to favour reinforcing incumbency rather than encourage innovation – i.e. the managerial rather than the entrepreneurial mindset – and for this reason are thus skewed toward a trajectory that favours probable over plausible worlds – i.e. circle 3 over circle 2 in Figure 1. Small wonder that a managerial mindset dominates managerial practice, and thereby knowledge management practice and even managerial education today.

Can we reduce the costs of firm-level creative destruction to the losers? Indeed, do they have to be cast as losers at all? When viewed from an evolutionary perspective, one of epistemology's key insights was that 'hypotheses could die in our stead'.²³ Embodied in appropriate policies, this insight should increase our tolerance for higher levels of uncertainty, and encourage us to aim for bolder hypotheses, hypotheses that strike us as plausible on the basis of their innate coherence rather than of their initial correspondence with the facts. To repeat: correspondence with the facts is typically not on offer in the early phases of a genuine innovation. The facts do not yet exist.

We can summarize the above points by means of Table 1, in which the cognitive inputs required by the different types of worlds that we have been discussing are each associated with specific kinds of knowledge management tools and where each creates a specific kind of value as an output.

World:	Knowledge Inputs	Actions	Outputs	Investment Criterion	Knowledge Management Processing Tools
Possibilities	Hunches	Real option investments	Real Option potential	Real option value	Brainstorming and scenario analysis
Plausibilities	Speculations	Speculative investments	Knightian rent potential (Potential entrepreneurial profits)	Subjective NPV	Pattern generation and recognition tools
Probabilities	Estimates	Risk adjusted investments	Profit potential (Potential normal profits)	Objective NPV	Statistical processing and analysis tools
Actualities	Calculations	Contracts	Profits (normal or abnormal)	Optimal solution	Optimization techniques

Table 1: Knowable Worlds and their Tools

Table 1 lays out in a systems format the challenges that confront the practice of knowledge management: inputs required, actions required, outputs expected and criteria for investment decisions. These will vary depending on the epistemological space you are operating in. We progress from possibilities towards actualities as we move down the table. Note that profits are an output only in an actual world; in all other worlds profits have potential, and with different degrees of likelihood.

The spaces in which knowledge management practices are currently the least developed, and therefore the least effective, are those describing possible worlds and plausible worlds - especially possible worlds, for which to our knowledge, few really successful practices have been created. Yet, in the absence of credible and powerful knowledge management methodologies that can identify hunches and evaluate the option potential of hunches, identify speculative investments and assess their Knightian rent potential, managers

will naturally gravitate to risk adjusted investments, with their more certain and therefore normal, as opposed to abnormal, profit potential.

Confining our epistemologies to a managerial mindset focussed on the worlds of probabilities and actualities in the bottom two rows of the table thus denies us access both to the rent potential of option values and access to Knightian rent potential.

In a world of pure possibilities, for example, we can invest in low cost real options based on hunches derived from vaguely linked beliefs, which if successful create real option potential that has option value. In this world, knowledge management can draw off and improve brainstorming and scenario generation processes with a view to helping the firm to surface more and better possibilities and assess them as option candidates. One reason why scenario planning has had difficulty establishing itself is because it has not articulated an explicit relationship between its epistemology and the process of creating of option values. Evaluation of such opportunities calls for knowledge management tools that identify those options opportunities with superior properties: upside potential , potential control of downside and potential sustainability.²⁴

On the plausibility front, knowledge management tools that enhance pattern generation and recognition need to be developed that will encourage searches for patterns of coherence and thereby more aggressive generation and exploration of plausible opportunities to commit resources to speculative investments that the firm alone can see. Much more aggressive development of text processing and text pattern recognition methodologies would be a major place to start.

In conclusion, by highlighting the extent to which knowledge management has been concentrating on the bottom two rows of Table 1, we can begin to appreciate both how far it

has been underselling itself as well as the scope that it has for broadening its agenda. Knowledge management can and should restore epistemic legitimacy to an entrepreneurial mindset that is primarily located in the top two rows of the table. This is partly a matter of cognitive style but partly also a matter of education. We keep teaching people that there are right and wrong answers to be derived by analysis. In some areas, to be sure, there will be; in others, however, we ourselves create the right answers by *enacting* them, thus converting a possible or a plausible prospect into a probable or an actual one. Reduce the social, organizational, and personal costs of operating in possible or in plausible worlds, improve the payoff for doing so, and more people will start behaving entrepreneurially. But to take on this task, knowledge management needs to come to grips with epistemological issues.

5. Conclusion

In this paper, we have made four points:

1. The Managerial and entrepreneurial mindsets operate with different epistemologies that reflect differences in their respective circumstances. The entrepreneurial mindset operates under conditions of novelty and uncertainty, ones in which prior probability distributions, being non-existent, can offer little guidance, and in which coherence is the epistemological underpinning. The managerial mindset, by contrast, is constrained to seek justification from probability distributions and thus cannot operate in their absence, so that correspondence is the epistemological underpinning.
2. Within each mindset the agent navigates from a world of possibilities to one of actualities, but through a different trajectory. The entrepreneurial mindset attempts to enact bold yet plausible hypotheses that create their own reality; the managerial mindset acts on the basis of objectively verifiable facts and constraints from which it derives hypotheses with a

high degree of probability. The first epistemic strategy will be appropriate under conditions of novelty and uncertainty. The second will be appropriate where accumulated and objectively verifiable experience is available.

4. Current institutional practice is heavily skewed in favour of the managerial mindset. Knowledge management could help to redress the balance, but only if it becomes more epistemologically aware.

5. In particular there is a need for knowledge management tools that allow agents to develop and explore possibilities –especially to generate and process possibilities- and also to develop and explore plausibilities -especially generate more and bolder ones. Seed tools already exist but need more epistemologically oriented development. These are brainstorming and scenario planning for possibilities, and pattern recognition tools like text processing and analysis for plausibilities.

The above has implications for theory. Much of modern management thinking has been inspired by the success of large established enterprises that operate on the basis of well-tested routines, well-documented facts, and hence articulable probability distributions. Starting with the Scientific Management movement at the beginning of the twentieth century, therefore, the managerial mindset has come to dominate both management education and management practice. As a result, the entrepreneurial mindset has got crowded out.

Yet the entrepreneurial mindset, it turns out, may have a broader scope than what is offered by purely entrepreneurial settings. As the level of uncertainty goes up in many walks of life, the entrepreneurial approach to the management of knowledge will surely grow in importance. When should we use it and when should we have recourse to a more traditional epistemological stance? Answers to this question have important implications for the way we

select governance mechanisms appropriate to different mixes of risk and uncertainty. In a regime of multiple and complementary epistemologies, a “one size fits all” approach to the problem of corporate governance cannot convincingly deliver. Yet, the near-total absence of any knowledge management contribution to the debates provoked by the recent Enron, WordCom, and other scandals, suggests that the discipline has not begun to address the challenges that it faces at the level of governance.

Our findings also have implication for Practice. We have stressed that we are dealing with mindsets and not with particular individuals. Large corporations are as much in need of entrepreneurial mindsets as they are of managerial ones. However, building an entrepreneurial mindset will require the explicit development of internal organizational procedures that both counterbalance and complement existing managerial mindsets. Corporate venturing activities, for example, have often suffered when submitted to internal evaluation procedures that reflect the predominance of the managerial mindset.

New procedures need to be put in place in organizations that aspire to entrepreneurial behaviour, procedures that acknowledge the contingent nature of the epistemologies under which organizations operate. The procedures will vary from firm to firm, but we believe that at a minimum they should observe the following principles:

- Develop systematic procedures for distinguishing between risk and uncertainty. Is the situation repetitive and hence amenable to probabilistic thinking, or is it so novel that little or no prior data can really guide a decision?
- Separate out the management of uncertainty from that of risk. They are not the same.

If a proposed venture is uncertain rather than risky – that is, neither the full range of

possible outcomes nor their respective likelihood can be identified - do not submit it to the quantification requirements of a risky situation.

- In order to manage uncertainty, develop skills in creating and linking together options that can be exercised in different sequences as the states of nature become known.
- Establish in which of the different worlds – possible, plausible, probable, actual – a venture proposal is currently located, and for each develop systematic evaluation criteria that are appropriate both to what can reasonably be known as well as to the chosen trajectory from one world to another. As suggested above, for example, do not ask for quantified estimates in possible or plausible worlds where none are to be had. In plausible worlds, focus on the coherence of a proposal and on the consistency of its underlying assumptions rather than on their likelihood.
- When dealing with entrepreneurial proposals, move away from the analytically oriented business plan questions of ‘how do you know that this will happen?’ This is the possible-to-probable world trajectory of managers. Instead, move toward the more action-oriented entrepreneurial question of ‘what actions will you undertake to make this happen?’ This is the possible-to-plausible world trajectory of entrepreneurs. In order to evaluate entrepreneurial proposals, insist on the construction of scenarios for the possible outcomes of such actions together with the identification of the options available should something like the scenario come to pass.
- Until there is a cadre of internal people with an entrepreneurial mindset, choose outside entrepreneurs to carry out the evaluation of plausible and possible proposals rather than managers.

- Avoid procedures that filter out entrepreneurial proposals prematurely. Allow such ideas to ferment during the evaluation phase. If they are radically new, they may need getting used to.

Appendix: Epistemological Issues

Since Plato, we have assumed that knowledge has a truth condition – i.e. a proposition either corresponds with the facts, is coherent with a wider system of propositions, or has pragmatic cognitive value. It can, of course have all of these. Also, knowledge implies belief but the converse is not generally held to be true.²⁵ A belief is a dispositional psychological state and a belief can be false. To satisfy the belief condition for knowledge, such satisfaction must be “appropriately related” to the satisfaction of its truth condition. A proposition is *epistemically permissible*²⁶ if consistent with a given set of epistemic rules of evidence; it is epistemically good if based on adequate grounds. Less normative, perhaps, is the demand that the justification of a proposition be based on evidential support. A contextualist view of justification that is endorsed, for example, by Dewey, Wittgenstein and Thomas Kuhn, holds that all justified beliefs depend for their evidential support on some unjustified beliefs that themselves need no justification. These will vary from context to context and from social group to social group²⁷. Yet whatever one’s approach, questions of justification attract the lion’s share of attention in contemporary epistemology.

Epistemology concerns itself with the limits of knowledge – and by implication with its scope. The more restrictively we draw the boundaries around what we take to be knowledge, the more skeptical we are. The most extreme forms of skepticism hold that we are not actually justified in believing anything at all!²⁸ But is the justification of beliefs to be

based on their predictive performance or on the understanding they give rise to? If on the first, then the quantum theory must be counted an outstandingly successful instance of justification – after all, the whole of modern electronics is based on it. If on the second, then it must be recognized that the theory remains highly problematic.²⁹ Feynman famously remarked that, whatever they might claim, no one really understands the quantum theory.³⁰

In recent decades, justification has been progressively “naturalized”.³¹ In contrast to theories of justified beliefs that concentrate on probabilistic or logical relations between a hypothesis and the evidence for it, the naturalized approach looks at the psychological processes at work in developing the belief. Alvin Goldman, for one, holds that one can only justify a belief if one can show that it has been produced by a reliable belief-forming process.³² But where, exactly, should we locate such a ‘belief-forming process’? Dretske (1981), for example, adopts an information-theoretic approach that has the possibilities for knowledge dependent on the physical capacity of instruments, gauges, neurons, etc, to pick up and process signals from the environment, extracting the relevant information from them. And almost as a natural extension of such an approach, naturalistic epistemology has been linked both to artificial intelligence and to the history of science.³³ Inspired by models of population biology and economics, naturalistic epistemology has also been analyzed from a social perspective in which both power and authority relations can shape beliefs.³⁴

Given these trends, we will not be surprised to discover that naturalistic epistemology had developed an affinity with evolutionary epistemology, a term coined by the social psychologist Donald Campbell to describe a theory of knowledge that is both inspired by and derived from organic evolutionary processes. One variant of such a theory that goes back to “Darwin’s bulldog”, T.H. Huxley, sees the growth of knowledge – particularly scientific

knowledge – as *analogous* to the transformation of organisms through processes of variation, selection and retention.³⁵ Another variant, which was espoused by Darwin himself, sees the process as *literal*, with our thinking capacities being channeled in certain directions by the biological selection processes that we are subject to.³⁶ Some supporters of the latter variant identify with Quine’s naturalistic epistemology. A naturalistic epistemology takes the human subject as a natural phenomenon and studies its epistemic activity empirically. Quine argued that epistemology should be a branch of psychology, a study of how organisms take sensory stimulations as inputs and use them to output theories of the three-dimensional world. It is descriptive rather than normative although it could be made normative according to Quine.

Naturalistic epistemology, recognizing that in the real world, certainty is rarely, if ever, on offer, has more pragmatic aims. In line with a proposal originally put forward by the philosopher of probability, F.P. Ramsey, it is happy to take as knowledge whatever subsets of an organism’s beliefs it is willing to act upon providing that these contribute to its survival and – in the case of human organisms – to its prosperity.³⁷ A requirement for justification and objectivity in this scheme merely translates into a requirement for inter-subjective agreement. A requirement for truth translates into a good fit between an agent’s beliefs and his or her stock of prior experiences rather than between such beliefs and some “God’s eye” view of the states of the world. As Lackoff and Putnam have each shown, such an approach to knowledge, by linking it to the conditions under which action is possible, opens up the field to multiple, non-exclusive epistemologies.³⁸ ³⁹ If the wilder forms of relativism are to be avoided, however, it poses the challenge of understanding the nature of the *boundaries* that might help one to distinguish one epistemology from the other, and by implication, the circumstances under which each can validly be drawn upon.

References

-
- ¹ Ashby, Ross W., *An Introduction to Cybernetics*, London: Methuen, 1956
- ² Boisot, M. *Knowledge Assets. Securing Competitive Advantage in the Information Economy* Oxford University Press, 1998
- ³ Clark, A., *Being There: Putting Brain, Body and World Together Again*, Cambridge, Mass: The MIT Press, 1997
- ⁴ Audi, Op cit.
- ⁵ Ayer, A. *The Problem of Knowledge*, London: MacMillan, 1956
- ⁶ Polanyi, M., *Personal Knowledge: Towards a Post-critical Philosophy*, London: Routledge and Kegan Paul, 1958
- ⁷ Thelen, E., and Smith, L., *A Dynamic Systems Approach to the Development of Cognition and Action*, Cambridge: Mass: MIT Press, 1994
- ⁸ Feyerabend, P., "Against Method" in *Minnesota Studies for the Philosophy of Science* 4, Minnesota: University of Minnesota Press, 1970
- ⁹ Borel, E., *Valeur Pratique et Philosophie des Probabilites*, Paris: Gauthier-Villars, 1937
- ¹⁰ James, W., (2000), *Pragmatism* (1907), New York: Penguin Books
- ¹¹ Castells, M., *The Rise of the Network Society*, Oxford: Blackwell Publishers, 1996
- ¹² Copp, N., and Zanella, A., *Discovery, Innovation, and Risk: Case Studies in Science and Technology*, Cambridge, Mass: The MIT Press, 1993
- ¹³ Derry, T., and Williams, T., *A Short History of Technology: From the Earliest Times to A.D. 1900*, Oxford: Oxford University Press, 1960
- ¹⁴ Cardwell, D., *The Fontana History of Technology*, London: Fontana Press, 1994
- ¹⁵ Derry and Williams, 1960, Op cit.
- ¹⁶ Ziman, J., *Public Knowledge: the Social Dimension of Science*, Cambridge: Cambridge University Press, 1968
- ¹⁷ Audi, R. *The Cambridge Dictionary of Philosophy*, Cambridge: Cambridge University Press, 1995
- ¹⁸ Thagard, P., *Coherence in Thought and Action*, Cambridge, Mass: MIT Press, 2000
- ¹⁹ Mary Douglas, 1966 *Purity and danger: An analysis of the concepts of pollution and taboo*, London: Routledge and Kegan Paul.
- ²⁰ R.Nisbett, 2003, *The Geography of Thought: How Asians and Westerners think differently... and why*, New York: The Free Press.
- ²¹ Holland, J., *Adaptation in Natural and Artificial System: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, Cambridge, Mass: MIT Press, 1975
- ²² Schumpeter, J., *The Theory of Economic Development: An Inquiry into Profits, Capital, Credit, Interest and the Business Cycle*, London: Oxford University Press, 1961 (1934),
- ²³ Popper, K.R., (1972) *Objective Knowledge: An Evolutionary Approach*, Oxford: The Clarendon Press
- ²⁴ McGrath, R.G. *A real options logic for initiating technology positioning investments*. Academy of Management Review. October, Volume 22, Number 4, pages 974-96.
- ²⁵ In this paper, we leave aside the issue raised by the so-called *Gettier problem*. In a highly influential paper published in 1963, Edmund Gettier challenged the equation between justified true belief and knowledge (Gettier, 1963). Given that Gettier raises issues that are both beyond the scope of our discussion and not particularly relevant to it, we shall take that equation as an anchoring point for our discussion.
- ²⁶ Chisolm, R., 'The Indispensability of Internal Justification', *Syntheses* 74, pp 285-96, 1988
- ²⁷ Kuhn, T., *The Structure of Scientific Revolutions*, Chicago: The University of Chicago Press, 1962
- ²⁸ Unger, P., 'A defense of Skepticism', *Philosophical Review*, 80, pp.198-219, 1971
- ²⁹ Miller, A., *Insights of Genius: Imagery and Creativity in Science and Art*, Cambridge, Mass: The MIT Press, 2000
- ³⁰ Omnes, R., Quantum Philosophy: Understanding and Interpreting Contemporary Science, Princeton: Princeton University Press, 1999

-
- ³¹ Quine, W., 'Naturalized Epistemology' in *Ontological Relativity and other Essays*, New York: Columbia University Press, pp. 69-90, 1969
- ³² Goldman, A., 'A Causal Theory of Knowing', *The Journal of Philosophy*, 64, pp.357-72, 1967
- ³³ Laudan, L., *Progress and Its Problems: Towards a theory of Scientific Growth*, Berkeley: University of California Press, 1977
- ³⁴ Hull, D., *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*, Chicago: The University of Chicago Press, 1988
- ³⁵ Campbell, D., 'Evolutionary Epistemology', in P. Schlipp (ed), *The Philosophy of Karl Popper*, LaSalle: Open Court, pp. 413-463, 1974
- ³⁶ Barkow, J., Cosmides, L., and Tooby, J., *The Adapted Mind: Evolutionary Psychology and the Generation of Culture*, Oxford: Oxford University Press, 1992
- ³⁷ Ramsey, F (1926) 'Truth and Probability' in Ramsey, F., (1931), *The Foundations of Mathematics and Other Logical Essays*, London: Routledge and Kegan Paul.
- ³⁸ Lackoff, G., (1987) *Women, Fire, and Dangerous Things: What Categories Reveal about the Mind*, Chicago: The University of Chicago Press,
- ³⁹ Putnam, H., (1975) *Mind, Language and Reality: Philosophical Papers, Vol. 2*, Cambridge: Cambridge University Press.