

Review Article: An Essay of the State of Knowledge Management

Architectures of Knowledge: Firms, Capabilities, and Communities

Ash Amin and Patrick Cohendet

Oxford, Oxford University Press, 2004, xv + 179 pp., UK£18, ISBN 0199253331 pbk

Organizational Knowledge in the Making: How Firms Create, Use, and Institutionalize Knowledge

Gerardo Patriotta

Oxford, Oxford University Press, 2003, 248 pp., UK£47.50, ISBN 0199256780 hbk

Organizations as Knowledge Systems: Knowledge, Learning and Dynamic Capabilities

Haridimos Tsoukas and Nikolaos Mylonopoulos (Eds)

Basingstoke, Hants, Palgrave Macmillan, 2004, xx + 307 pp., UK£50, ISBN 1-4039-1140-1 hbk

Some Background Issues around Knowledge Management

Anyone who looks into 'knowledge management' (KM) sees immediately that the term knowledge has been co-opted to a wide variety of agendas—many commercially-driven and likely no part of KM. We have neologisms such as the Knowledge Age, knowledge workers, and Knowledge-Intensive Firms, as well as a systematic re-labeling of activities as 'knowledge-based'. But what is really new here? Since when did our theorizing not make data or information central? Where can we find the difference between administrative and production rules and organizational 'knowledge'? Clearly much of today's KM consulting is old-style systems analysis, human resource management, or organization development, updated and re-labeled. There is a great deal of hype, not only in the trade press, and an already apparent danger, that over-use of the K-label is going to shut down scholarly and managerial interest in the novel thinking that these books, and the KM field generally, often advances. We need some framing, some clarity on what, precisely, KM is about—other than the mere opposition between knowledge and its absence.

The bulk of the managerial excitement, and the core of the more accessible KM literature, approaches it as the better identification and management of the organization's 'knowledge assets', where that term implies some non-traditional type of resource. The term 'tacit' is used to point to employees' knowledge and skills which, though not on the balance sheet, are clearly both expensive and important to the organization's operations and competitive potential. Customer and supplier relations, and goodwill, can be similarly considered. In this way KM seeks to extend our traditional notions of decision-making about the acquisition and allocation of tangible assets to cover intangible assets as well. The implication is that organizations contain or possess under-considered assets which may be strategically important, especially if they are the source of sustainable economic rents.

Human capital accounting, as popularized by Skandia, is part of KM in that it attempts to extend conventional accounting to lead to better management of such assets. Likewise much of the organizational learning literature is focused on their creation, extending product and process development to embrace the development of new K-assets. There is also a communication angle. If K-assets are information-like we can wonder about their mobility and ask 'is conventional IT able to meet the challenge of moving these around in a timely manner to where they are most needed?' We are all familiar with the challenges of getting new R&D into production. Much of the KM literature deals with using modern computing to correct such mal-distribution. Terms like 'sticky' and 'slippery' are widely used. Between organizations there is an associated attention to 'absorptive capacity'. Theory about such matters would be a proper part of KM, extending conventional organizational communication theorizing.

In general KM has grown out of our interest in some different types of knowledge asset, such as employee skill or an organization's dynamic capability, and ways of bringing that asset under current managerial theory. A good example is Teece's work on managing intellectual capital. *Inter alia*, he argues a key problem with knowledge is its 'public good' or non-rivalrous nature. Effective intellectual property rights 'secure' this knowledge and render it rivalrous, at which point it can be priced and managed in the same manner as the firm's other rivalrous resources. Likewise Nonaka and Takeuchi's work on managing the conversion of the employees' tacit knowledge into the firm's explicit knowledge performs a similar service. Boisot's work on codification, abstraction, and diffusion addresses the firm's knowledge identification and movement challenges. Learning organizations and communities of practice are similarly seen as ways of generating organizational knowledge that go beyond rationally planned R&D.

This kind of KM is about identifying as wide a range of knowledge assets as possible—be they forms of knowledge, or knowing, or proficient practice at the individual, group, organizational, cluster, industry, region, or national level—and bringing them into our theorizing about maximizing efficiency, profit, or market power. It is highly relevant to managers because they see these under-considered assets may have important strategic implications. Organizations normally generate slack resources, and the tendency to re-invent the wheel rather than seek out the local expert who has already done this is endemic. Better knowledge management can reduce this possibility. So KM's new resource locating and managing implications are of interest whether the objective is waste-minimizing or rent-maximizing, i.e. independent of the underlying theory of the firm.

Alternative Concepts of Knowledge Management

At this point the question must be: 'is this all there is to KM?' To answer 'yes' is to consign KM to the subordinate role of extending existing theories by helping them embrace intangible assets. Being able to account for human capital, or establish the transaction costs of identifying and relocating the firm's knowledge assets, or make economic decisions about organizational learning, seems to extend micro economics and organization theory without changing them greatly.

But it is at this point that the academic and practitioner interests begin to diverge. Practitioners will be happy to focus on managing these newly revealed resources. But academics will begin to look for concepts that bind the different theory extensions together, for that way lies a knowledge-based theory of the firm, a potentially significant advance beyond the asset-based theory. We can ask 'does human capital have some conceptual relationship with non-rivalrous resources?' Or 'Can organizational learning provide insights into how micro economics should develop?' Or 'Can a rule-based organization manage the employees' skills effectively?' In general, we can wonder whether KM provides a context in which traditional managerial theorizing can be challenged and perhaps transcended. It is the old question of how best to balance evidence and theory. The first strategy in our literature is to contain and resolve the curious and problematic aspects of knowledge, such as its tacit or non-rivalrous-ness, so that it can be brought into theories which have no space for such characteristics. A second strategy follows falsificationists such as Popper and Lakatos. When theory cannot handle the phenomena, it must be abandoned.

This takes us towards a radically different type of KM that uses the under-specified notion of knowledge to problematize one or other of the taken-for-granted concepts of managerial theorizing. The first strategy glosses over the problems of defining knowledge, ultimately hanging from some 'you'll know it when you see it' assumptions. It thereby implies rather than demonstrates that knowledge can be brought into mainstream theory. Our cautions about this give rise to frequent appeals for empirical research. The second strategy is the search for new theory, not accommodation within the old, and works only because it presumes we do not have a secure definition of knowledge. For example, to see knowledge as a corporate asset is to problematize the previously secure notion of corporate resources. If the firm's managers cannot effectively identify and secure that knowledge it cannot be drawn into their canonical decision-making models. In the face of such an attack—clearly implicit in Simon's notion of 'bounded rationality'—our theory of managing and of the firm/organization may need to be changed beyond the limits that its present certainty- and rationality-oriented axioms allow. This is KM's radical possibility. It is, of course, ironical that the most immediate message of this strategy is that KM might be better styled 'knowledge problem management'.

One widely adopted strategy for dealing with the manager's intractable knowledge problems is to complement her/him as a knowledge consuming decision-maker with some other knowing agent or system. For instance, the assumption behind the communities of practice literature is that such communities evolve their knowledge from practice in ways that cannot be explicitly modeled and so made theoretically or explicitly available to managers. Communities of practice are not knowledge-generating engines that can be readily harnessed to corporate goals, rather they are seen as well-springs of newly emergent knowledge and practice. Some writers assume the firm can be considered a single community of practice,

others that it might be better seen as a community of communities of practice, others that the products of communal practice evolve into identifiable organizational routines. The key here is to see 'community practice' as axiomatically distinct from collective rational decision-making or discovery. As a result, it radicalizes our theory of organization, which is forced to move away from conventional axioms of methodological individualism and unity of purpose and command, and into a more complex framework of multiple and probably incommensurate actors. The same intellectual maneuver has been worked at the network, cluster, region, and industry levels. Each explanation then relies on some under-specified notions of 'learning-by-doing' or of how collective practice leads to knowledge, thereby radicalizing our conventional theories of management. Practice is no longer just the goal-oriented implementation of a manager's rational decision, rather it is treated as a concept on its own, cut free from the manager's rationalizing, an independent and possibly primary source of the knowledge that managers lack.

To suggest that such 'knowledge-based' theories are testable or even useful is to get way ahead of the literature. But the direction of such progress is becoming clearer and, frankly, things are looking good. These conjectures could shift managerial theorizing into some radically new space, sometimes alongside some other cognizing agent, such as a community of practice, sometimes to within the system with the presupposition that it is self-managing and/or 'autopoietic'. The methodological implications of these shifts are profound, so much so that we can see the first kind of KM as a jaded prisoner of the current positivistic hegemony, while the second bristles with more up-to-date possibilities—and difficult challenges. Post-modernism, radical constructivism, critical realism, and other terms fill the conversation.

But behind this lies something even more exciting. For decades our academy has been riven with concern about the seeming irrelevance of our theorizing, dividing us into at least two camps. One argues we need to do more of what we have been doing, only more rigorously; the other, perhaps most visibly led by Mintzberg, argues our difficulties follow directly from inappropriate methodological commitments. The accelerating use of 'knowledge' to problematize conventional theorizing, first taking out over-simple notions of 'resource', the key to Penrose's theory of the growth of the firm, and then rational decision-making, the key to the communities of practice literature, and then, perhaps, the organization itself, so opening us up to social capital and networking theory, drives us towards a major re-evaluation of our research methods. Given that we are unlikely ever to find good empirical support for including such aggressive notions as organizational routines or dynamic capabilities within our established theories, we may need to explore those research methodologies that gain us better insight into systems of social practice among multiple and disparately creative actants.

A second conceptual maneuver in the radical KM literature is to follow authors such as Weick and Orlikowski and argue for 'knowing' rather than 'knowledge', so bringing the organization's on-going and recursive practices center-stage. The focus shifts from asset-like knowledge objects, separable from those generating and using them, towards situated systems of social practice. It is not always clear, however, whether this maneuver is simply the first maneuver re-labeled, an appeal to alternative sources of knowledge, or a move to a different epistemology or form of knowledge. Hidden here is the century-old social science struggle between objectivity and subjectivity, now more usually framed as that between the realist/positivist and interpretive epistemologies. While the former sees knowledge as about a

reality 'out there', the latter sees what is known as a mere interpretation of sense-experience. In the latter epistemology there are no correspondence-rules against which knowledge can be checked, empirically or otherwise, so truth criteria are likely to be socially constructed. Facts derive more from the social processes of generating and warranting knowledge than from whether the result represents 'reality'.

The shift to an interpretive epistemology does not necessarily force us to abandon the first type of KM. As Kuhn showed, within a stable language-game or knowledge paradigm interpretations take on the properties of established facts, widely taken for granted. Popper and others taught us to leave the search for ultimate proof of knowledge's correctness in the wings, for to bring it on stage is to embrace solipsism, to deny the concept of knowledge. More revealing, once we see that anarchic relativism can similarly destroy the intellectual enterprise, is the contrast between the pragmatic and social practice approaches to warranting knowledge. The former is crude and to the point, look at what works within the scope of a particular event. Here questions about knowledge being contained in language and/or communicated are put aside. In contrast, social warranting depends crucially on language and on whether statements can be understood and judged correct by others. In practical terms the community of practice is not simply phenomenological, it becomes manifest in informative discourse. To use the fundamental insecurity of knowledge as the gouge to radicalize managerial theorizing we must deal with the realist defense that, given further research all ignorance will disappear. This presents the theorist with some profound methodological challenges and the second KM strategy stands squarely on dealing convincingly with these. Merely invoking the term 'knowing' does not do the trick. Conversely, there can be no radicalization of managerial theorizing without a related radicalizing of our research methods, for the two are completely intertwined. In short, we see KM could change our most fundamental concepts of management—a reason for excitement as well as trepidation.

To the Books Themselves

These three books explore this new and radical agenda in revealingly different ways. The most overtly aggressive is the conference-based essay collection edited by Tsoukas and Mylonopoulos. This has at least four conceptual pieces that deserve to be widely read and cited. Patriotta's book is equally aggressive, though more narrowly in its search for a workable KM research methodology that can explore how organizational knowledge is developed and applied in a real case situation. More conventional and even-handed is Amin and Cohendet's exploration of collective knowledge generation, which follows their masterly review of the KM literature.

Amin and Cohendet: *Architectures of Knowledge*

Amin is a socio-spatial geographer at Durham University in the UK, Cohendet a public-sector economist at the University of Strasbourg in France. Both are senior members of their professions, widely published, with working relationships to some of the 'big names' in our field as well as experience of the practicalities of European social science research and policy-making in Brussels. They come to KM with a wider bibliography than the majority of us already in the field, though their grasp of our literature is firm. Their objective is revealed in their subtitle: firms, capabili-

ties, and communities—to understand how social communities develop economically significant knowledge capabilities, and how these become central to firms. The agenda is one a number of other authors have explored, but they also deal with seldom noted issues such as corporate governance and personnel incentives, and the public policy questions raised by a knowledge based understanding of how firms work.

Their thesis is that the ‘true spark of innovation lies in the generative dance between possessed and practised knowledge’ (p. xiv), which they take from Brown and Duguid. The interplay between ‘knowledge’ (knowledge-as-asset) and ‘knowing’ (knowledge-as-practiced) takes place in communities sharing similar practices. Chapter 1 reviews the various notions of knowledge. Three K-based theories of organization emerge: (a) a top-down ‘strategic management’ approach emphasizing knowledge application and management by design; (b) an evolutionary-economics approach arising from evolving organizational routines; and (c) a social anthropology of learning focused on the working community as the knowing and learning actant. Chapter 2 reviews the economic approach to knowledge, questioning in turn whether knowledge can be (a) increased by simple accumulation; (b) invariably codified; (c) held only by individuals; or (d) considered an asset ‘possessed’. This sets up their argument that knowledge, and rationality, is essentially partial and dynamic and thus an aspect of a social process that, first of all, needs a place. Chapter 3 proposes the firm as this space, a real-world locus for competence building. The firm as a decision-making and transaction managing apparatus is contrasted with the firm as ‘an institution where competences are continuously built, shaped, maintained, and protected’. The outcomes are organizational routines, the knowledge products of an evolutionary learning process and the local governance mechanisms for those engaged in that process. But the key, they argue, is the ‘intermediate level between the analysis of the behavior of individuals and the behavior of the whole organization’. This is the place of ‘communities’.

Chapter 4 is an analysis of the collective processes that lead to learning. Much of the chapter compares and contrasts the ‘communities of practice’ and ‘epistemic communities’ literatures, with a discussion of actor-network theory thrown in for good measure. Their conclusion is that KM fashion has swung too rashly from overly cognitive models towards overly socialized notions of social capital and communal trust, thereby overlooking the real and tentative fabric of the relations which both constitute community and are the channels of its learning processes. This sets the stage for Chapter 5’s analysis of spatial arrangements and the impact they have on these relations. The ‘regions of innovative activity’ arguments, illustrated by Silicon Valley, Route 128, and Emilia Romagna, are well known and go back to Marshall and beyond. But the authors tease out important subtleties concerned with the circulation of knowledge as the essential prerequisite to learning, proximity alone being insufficient. They remind us that these spatial relations and interactions need to be managed, not left to chance meetings at the watering holes of Palo Alto.

Chapter 6 addresses the management of learning communities directly, contrasting ‘management by design’ with ‘management by community’. In the spirit of March’s balancing ‘exploration’ with ‘exploitation’, the authors seek a balance between sustaining and supporting the existing practices and the exploration of novel ones. Finally in Chapter 7, with a rich and dynamic image of community-embedded knowledge generation in place, they analyze the debates

underpinning national R&D programs and associated institutional developments. They remind us, for instance, that the patent system does not adequately motivate the complex international private-sector networks we now see driving innovation. Their recommendations are both practical and useful calling, for instance, for a better match between the managerial, legislative, and institutional systems and the real corporate-driven innovation processes now evident throughout the global economy. This slim and well-produced book shows we must now move far beyond narrow discussions of managing communities of practice and learning organizations and consider the broader implications for a changing socio-economy increasingly driven by knowledge.

Patriotta: *Organizational Knowledge in the Making*

Patriotta's book, also blessedly slim, also pivots around collective learning. At its core is a case study of the early 1990s building and operation of Fiat's new factory at Melfi, east of Naples in Southern Italy. The outcome was remarkable, creating one of the world's most technologically advanced and successful assembly plants. This case is contrasted with another about Fiat's older body pressing plant at Mirafiori near Turin, the scene of many management-union struggles. On the face of it, the story is one of establishing a green-field non-union site—but that misses the point. As the contrast between the US Big-3's Southern plants and the Japanese and German US transplants shows, the absence of union-based impediments does not, on its own, ensure success—it merely changes the context of the challenges. Patriotta's story is ultimately about the implementation of a business model that broke new ground beyond Japanese implementations of 'lean production' and is now dubbed the 'integrated factory'. It entails aggressively novel ideas about the relationships between leading edge technology and empowered staff and, pivotally, about how to make good use of collective learning.

But most of all, Patriotta gives us new insight into how to do empirical research in a KM frame. Ironically we see the researchers' concerns must begin to converge on the managers', as both focus on discovering the processes of knowledge creation. The result is not knowledge that then serves the traditional design and control paradigm. On the contrary, its power lies in emancipating managers from their taken-for-granted practices and moving them towards an acceptance of their employees' ability to resolve certain types of problems and thus a new division of labor. The dust-jacket blurb tells us: 'this book does not provide any recipes about alleged best ways for managing organizational knowledge. Rather, it invites managers and practitioners to reflect about the repertoire of knowledge they possess and yet cannot articulate' and so becomes a practical illustration of how to approach the methodological challenges mentioned earlier.

The book is in three parts; the first dealing with epistemological foundations, the second the case studies themselves, the final speculates on a knowledge-based theory of organizations. The introduction (Chapter 1) opens with a discussion of Milan Kundera's observation of modernity as noise, how this has changed the nature of life and being, reversing foreground and background. Patriotta's point is that the world presents itself interpreted in terms of unstable relationships between the 'subsidiary' taken-for-granted and the 'focal'. It sets up his discussion of method which is central to this book's achievement. He observes that the close analysis of everyday practices in organizations seems to pose profound challenges to mainstream theories of knowledge—and management. He questions (a)

whether knowledge can be commodified and managed; (b) the causal links between knowledge and performance; (c) the transience of organizational knowledge; (d) the alienation of business theory from managerial practice; and (e) the privileging of knowledge-creation over the dynamics of institutionalizing—or appropriating the value of—knowledge. At this point Patriotta's debt to Pettigrew, his advisor and the author *inter alia* of a magisterial case study of managerial processes in the UK's giant Imperial Chemical Industries, becomes fully clear. This work is about going deep into the practicalities of knowledge management in huge enterprises; no loose waffle about knowledge workers here.

Chapter 2 critically reviews existing perspectives on knowledge in organizations with the intent of identifying gaps. Patriotta suggests four paradigms: knowledge as representation, knowledge as commodity, the situated approach, and knowledge and the laboratory, such as Latour's work. He concludes that an interpretive approach that yields representations of reality separates our inner and outer worlds and thus overlooks the collective public nature of much important knowledge. The commodification approach sees knowledge in overly functional terms and so side-steps the problems that organizations have with trying to manage its non-zero-cost creation. While the situated approach has a strong empirical tradition, it seems to eschew generalization and thus conventional theorizing. The laboratory or technoscience approach stresses the processes of social construction but also risks relativism. While much more could be made of these contrasts, Patriotta is really justifying his adoption of a pluralist approach, one that admits all these paradigms, and thereby taking up Geertz's 'thick description' as the core of his research method.

Chapter 3 fleshes out this method and it is excellent. The key is Patriotta's development of three working concepts: time, breakdown, and narrative. Time is almost completely overlooked in the KM literature, yet is clearly crucial to understanding how knowledge moves from the point or locus of its development to that of its application. Patriotta surmises that knowledge is developed collectively and must be institutionalized before it becomes available, aligning him somewhat with Nonaka and Takeuchi's work. While this process can be observed, it can only be regarded in terms of the parallel flows of time and knowledge, a Western approach quite different from Nonaka and Takeuchi's attitude to time. Breakdowns call the taken-for-granted into question, so de-institutionalizing the organization's knowledge and shifting it from the background of the domain of practice into the foreground domain of analysis and discourse. Finally the study of narrative reveals the knowledge carried therein, so setting up a dynamic and creative dialectic between the collective practices that generate and apply organizational knowledge and the discourse which both shapes and is shaped by it. Patriotta believes that the results of applying these research methods can be made available by using Geertz's techniques to weave 'webs of signification'. Though the outcome is scarcely recognizable or legitimate to those of a realist disposition, it is sophisticated qualitative research and takes us well beyond 'mental maps' as a methodology for the interpretive position.

The case section begins with Chapter 4. The history of Fiat's struggle to find manageable and acceptably efficient processes of auto assembly sets the context in which the organizational knowledge is situated. Fiat was a committed adopter of Fordist technologies, yet among the first in the world to turn to wide-scale automation. Between 1980 and 1986 they reduced the workforce by 42%. GM's Roger Smith was also embarking on his total-automation killer-strategy, designed to solve his productivity crisis with capital expenditures that neither Ford nor Chrysler

could afford. Unlike GM, which stumbled deep into disaster, Fiat found good returns to automation, doubling their labor productivity and cutting the total labor charge from 28% to 17% of turnover. But it proved no defense against the Japanese auto firms' attention to quality. At the end of 1989 Fiat's CEO moved to develop a different strategy, their own approach to 'lean production' that was to become the 'integrated factory'.

While no less technology-intensive the objective of automation shifts from deskilling and substituting for labor and becomes a tool in the hands of an empowered workforce. The cultural and organizational implications of this shift are legion, of course, and Patriotta touches on these without losing sight of his objective, which is to look at the KM aspects rather than the industrial sociology or competitive strategy. But he misses an opportunity to tell us his views on the relationship between the methodology chosen and the subject matter. We begin to see an architectural relationship between the rational quantitative methods so widely adopted in business research and the top-down rational/mechanical models of management and organization which Fiat, and the rest of the world, are now moving away from. The integrated factory is a pluralistic multi-nodal network of operating units, each with significant responsibility for its own existence, relationships, and practices. These are no longer determined centrally. Thick descriptions are expressly designed to penetrate and reveal such complex and contested domains, each reality respected as a participant in the collective discourse which both shapes and is shaped by the participants' struggles for identity. The implication is that we cannot investigate the kinds of business problems that interest Patriotta—and Fiat's senior managers—without abandoning the totalizing methods that rule our leading research journals.

Chapter 5, which tells the story of the construction of the Melfi factory, is concise but wonderfully full of detail about how Fiat set about doing something radically new. In short, they recruited with extreme care, took the people hired into an extensive training program so that they knew what Fiat and auto manufacturing was about, and then let them build their own factory. Fiat's bosses challenged the new employees to learn by doing, and then take responsibility for doing it for themselves. The title to the book's Section 5.4 'Learning as appropriation: building identity through ownership processes' says it perfectly. What, the cynical reader might think, of the emasculation of the unions by locating in a non-union location, what freedom do these new employees really have to disagree with the bosses in Turin? Patriotta deals well with these issues, such as those Best revealed were crucial to the evolution of the Emilia-Romagna industrial successes, and still keeps his focus on the main event, the bosses challenge to the work force framed in terms of current adult learning theory. Here Patriotta reveals the shallowness of much of the published material on organizational learning and its tendency to ignore the work of the educational theorists and developmental psychologists.

Likewise Chapter 6 richly details the process of diagnosing and so learning from breakdowns and bottlenecks. The focus is not on tacit knowledge, skills, or other vaguenesses, rather it is on the well-defined interface between the employees and the equipment they use to reach objectives they have set. The technology is treated as a tool, sometimes better than those more ready to hand, but not always. The Melfi personnel did what shop floor people always do, trusting their ears and eyes, rather than their instrumentation, to 'understand in real time what's going on'. Again there is a temptation to re-interpret Patriotta's interpretations, and that is the whole point of his method, webs of signification are inherently flexible,

multi-faceted, and open to 'interpretive ambiguity'. That said, he probably makes too little of the difference between Fiat's normal production model, in which breakdowns are analyzed by production engineers acting on behalf of the plant's designers, thus disconnected from the operatives, and that in the Melfi plant where the response capability is built into the production system itself by engaging the operators in the design of the plant.

Chapter 7 moves us to the Mirafiori Pressing Plant. At the same time Patriotta abandons 'time' and 'breakdown' as analytic tools and takes up 'narrative'. This is so smoothly done that it is only later that one realizes it deserves a book in itself, for Patriotta just implies that the Mirafiori experience is so unlike that at Melfi that different tools must be brought to bear. At Melfi his time instrument is able to observe the processes of knowledge generation and sedimentation (institutionalization). Likewise he sees dealing with breakdowns as a process for generating new knowledge and expunging the old. The implication is that these processes have long since ossified at Mirafiori. The operators' experience is thus not at all about learning and reshaping the socio-technical system in which they are embedded. Now all the operatives can do is manage the experience of operating an unchanging system which, since it no longer has the capacity to absorb the uncertainties that inevitably arise, leaves them free-floating, to be assigned by political game of the plant's 'blame system'. The narrative tool becomes a detective story; we need to see who could have prevented the crime as well as finding and punishing the culprit. This is so on-target it is creepy to anyone with this kind of organizational experience, and speaks volumes about the consequences of managers abandoning their responsibilities towards their organization's knowledge processes. Patriotta is then able to make points glossed earlier about the emotional dimensions of generating and applying knowledge. Feelings run high at both plants, but in quite different directions and with quite different consequences. Narrative, he tells us, supports the flow of information between people rather than machines. Time and breakdown are more about the man/machine system interface. Again we sense the importance of matching research method to subject matter.

The two final chapters (Chapters 8 and 9) are the sort of obligatory summing up that reminds one of reporting in class what one did during the vacation—very different from the experience. Patriotta struggles to build a framework that stands up against Nonaka's familiar knowledge-spiral. But in the process his time, breakdown, and narrative typology somehow transitions into blueprints, routines, and common-sense, while at the same time alluding to a typology explored by Bruner and others which distinguishes foundational knowledge from procedural and experiential knowledge. This does not seem to work too well or lead to new insights, perhaps because Patriotta is more of a down-from-the-verandah anthropologist than philosopher, which is clearly what is needed at this point.

There are interesting speculations such as the idea that knowledge-making is based on entropy, the presence of imbalance and discontinuity. He tells us his three research instruments—time, breakdown, and narrative—are informed by a unifying principle of discontinuities. But the philosophical implications of this are neither explained nor explored. He suggests: 'we are left without a clear definition of what organizational knowledge is' and argues that at best it is generated by controversy. It is probably about how, resolving controversies, organizations come to know themselves, make sense of their performance, and what they do on the basis of that understanding. Yet he says this without, it seems, being completely aware of his increasingly reflexive stance. Finally he asks what a Fiat manager might

learn from the work—and provides the blurb which appears on the book's dust-jacket. He thereby shies away from the deeper philosophical questions which follow from contrasting realist notions of knowledge as about the world against alternatives such as Marx's suggestion that it is about changing that world, the more pragmatic one of simply coping with the world, or the more reflexive notion that one's knowledge of the world is no more than a reflection of what one has done to construct it.

Tsoukas and Mylonopoulos: *Organizations as Systems of Knowledge*

Amin and Cohendet use the notion of knowledge to synthesize, drawing from fields they know well and so broadening the KM discourse to include human geography and public policy. Patriotta privileges observing knowledge as practice over theorizing, so getting himself closer to a group engaged in building both a new plant and a new way of working. Both books see collective learning as the key to knowledge growth and a sketched theory of the organization. Though neither book fully engages the philosophical issues raised by collective practice, they are important because they take us confidently beyond the bounds of contemporary theorizing as a result of proposing knowledge making processes that lie beyond and thus complement rational decision-making.

The Tsoukas and Mylonopoulos collection is overtly philosophical. It is, for the most part, the proceedings of the 3rd European Conference on Organizational Knowledge, Learning, and Capabilities hosted by ALBA (Athens Laboratory of Business Administration) in April 2002. Careful positioning by the organizers led contributors to focus on practice-based or constructive approaches to the firm/organization as a knowledge system. The search for an understanding of practice is the key to this volume's coherence and contribution. John Seely Brown's short foreword on the importance of respect reaches deeper than might appear on first reading. His conclusion, that inter-departmental relations will be more productive if departments treat each other with respect, seems mundane. But Seely Brown is not only an experienced manager. He gives us the subtle notion of 'listening with humility' and thereby stresses that inter-departmental relations go better when they are created collaboratively—not quite what most people would get out of the mundane advice. While there are echoes of Patriotta's story, and Weick's notion of mindfulness, the crux is that the practice of respectful collaboration reconstructs concepts, language, and perception in ways that go far beyond the conventional ideas of facilitating communication between different work-groups.

Three notable chapters—the Editors' introduction, Gourlay's on tacit knowledge, and Lewin and Massini's on innovating and imitating firms—explore the difficulties facing anyone wanting to theorize seriously about practice's relationship to knowledge. Lewin and Massini struggle with the differing practices of imitating and innovating firms, and the resulting relationship between these and the economically framed theory of the firm. The fact that they ultimately fail to justify such differences is beside the point, for their discussion of organizational capabilities is certainly one of the best available. Readers with little more than a fleeting acquaintance with Polanyi's writings will find Gourlay's chapter on tacit knowledge a gift. This term has been used to excess and, as a result, has as many meanings as yesteryear's 'paradigm'. Gourlay, drawing on Collins, considers three current meanings: (a) unarticulated motor skills, such as bicycle riding; (b) the rules-regress idea, reflecting our ultimate inability to justify the rules of action with

anything other than an appeal to institutionalized practices; and (c) a 'forms of life' approach. Gourlay notes the last may be confusion between the rule-regress idea and the associated Wittgensteinian proposition that a system's fundamental properties cannot be observed from within. He reviews work by Baumard, Boisot, and Boiral for whom the meaning of tacit is tied directly to the practice of codification; their differences being over whether or how the tacit can be codified.

In the face of this evident confusion Gourlay goes back to Polanyi and points out that he seldom spoke of tacit knowledge. He spoke more of tacit knowing; in particular implying a relationship between tacit knowing and the Gestalt notion of coherence, by definition beyond the reach of language. Gourlay's argument is somewhat analogous to Hume's skepticism about causality. Neither coherence nor causality are proven aspects of the world. We impose such characteristics on our sense data. For Gourlay, tacit knowing is the practice of creating the coherence exhibited by our law-like knowledge of the observed world. These laws are our constructions and are thus never of the world revealed in sense-data, so the tacit is obliged to stand outside the domain of sense data. Under these circumstances the phrase 'knowing more than we can say' has less to do with the codification of sense-data than with the limits of language and the resulting gap between language and experience, between representation and practice.

Gourlay goes on to consider semiotics and awareness, drawing on Dewey and Bentley, suggesting tacit knowing may be communicable through the non-verbal signing that goes on as novices watch and listen to experts at work, much as the Melfi operatives listen to their machines. This is an interesting angle, but it is not clear that this part of his chapter works as well as the earlier part. But Gourlay's idea of the 'tacit dimension' as the practice of selecting what to attend to and to imposing conceptual coherence upon it is especially useful. We see immediately this practice can never be separated from the entity 'doing' the practice. It is part of the process or practice of that entity's consciousness and identity. Thus the competent bike rider selectively attends to a number of sensations and actions which, when functionally relevant, take on coherence and drop into subsidiary awareness as proficiency. It thereby becomes part of her/his identity as a skilled practitioner in a practice-defined world. Such philosophizing moves us towards an epistemology of practice along the lines explored by radical constructivists such as Pickering,¹ referenced and discussed only in Amin and Cohendet.

Perhaps the most important chapter is the Editors' introduction. They make at least three notable points. First, they consider the shift since ancient times from a focus on self-knowledge to today's utilitarian other-oriented knowledge of the external world. They conclude that, as a result, what we now take to be factual knowledge has become abstracted and separated from its previously reflexive, situated, and evaluative dimensions. Consequently there is a Rylean gap between the ancients' way of knowing, entwined with living and being, their everyday practical common-sense, and today's purely cognitive knowledge about the world in which we live, representations abstracted from our everyday practice which thus inevitably miss something crucial.

So their second point is the fundamental phenomenological one that a phenomenon is always more than the sum of its manifold representations. The key loss is insight into the phenomenon's 'potential', what it might become through practice, for representations do not give us an understanding of practice. Practice, and the knowledge manifest in proficient practice, is unlike any representation of the world. While there are links to Ryle and Polanyi here, there are interestingly

differences. Tsoukas and Mylonopoulos use the distinction between representation and potential to problematize knowledge and, as their third point, make a place for managerial and organizational creativity. The gap between representation and situated social practice is their locus of innovation.

The volume's fourth stand-out chapter is Pitelis' on Penrose. There is a growing cottage industry explaining 'what Penrose really said', and Pitelis is a leading commentator along with Foss and Loasby. Their work is important because, as with Polanyi, too many with a fleeting grasp of Penrose's work suggest an association with the RBV. In contrast Pitelis argues that Penrose went a long way to restoring the balance between the older economic concern with wealth creation and the current neo-classical focus on distribution from which the RBV emerges. Rather than build her theory on maximizing, Penrose presumed creative agency, something beyond the limits of the RBV. This and the managerial team's evolved knowledge, are the keys to her theory. Pitelis suggests that firm-like creative collaborative practice is as fundamental to a socio-economy based on knowledge as is Hayek's argument that markets provide the optimal means for distributing such knowledge. Pitelis makes an intriguing point about Penrose's inattention to organizational conflict, suggesting that she believed organizational stability was a co-product of the practices that lead to innovation. Patriotta makes a similar point about how conflict among the Melfi personnel was contained by their shared commitment to learning.

The remaining chapters, with the exception of Argyris's on double-loop learning, are readable and workmanlike explorations of what it means to pay proper attention to practice when theorizing about organizations. Bogenreider and Noteboom explore the emergence of shared beliefs within work communities. Huysman critically re-examines the literature on communities of practice. Haeffliger and von Krogh add further to our understanding of how the open source software community's practices become so productive. Boer, van Baaalen and Kumar reconsider the practices of knowledge sharing and offer a novel typology of social relationships. Treleaven opens up the typically overlooked issue of power in knowledge-sharing relationships with a Foucauldian critical discourse case about creating a new school in an Australian university. Daskalaki and Blair explore activity theory as a basis for their research into the film-industry's semi-permanent workgroups and their knowledge generation, retention, and application processes.

The last two chapters by Soo, Devinney and Midgley, and Anyfioti, Dutta and Eveniou are interesting explorations of non-radical KM. While the first of these is inherently tautological, measuring 'knowledge quality' in terms of its impact on the firm's performance, the findings on the importance of informal networks for knowledge acquisition are provocative. Anyfioti *et al.*'s final chapter, dealing with customer relationship management (CRM), or one-on-one marketing, explores how, in specific contexts, firms make use of the ways in which customers define products in terms other than price.

A standout for quite different reasons is Argyris's opening chapter on double loop learning. This is remarkable for its lack of fit with the practice-oriented and Wittgensteinian thrust of the rest of the book. Learning, we are told, is the detection and correction of error. This takes us back to Plato's essences and presupposes actionable knowledge about what we clearly do not know, so it runs completely against what Polanyi, Penrose, and any social constructionist would tell us. Action, we are told, is 'behavior with meaning', the latter being the actor's intention—so there is nothing tacit here—and 'in order to produce effective action, it is necessary to begin with knowledge that is generalizable'. Thus Argyris sees knowledge itself as

unproblematic save in its absence. Practice is knowledge's implementation and thus not a concept with an axiom-shaping place in the theorizing. This is precisely the kind of explicit cognition-bound theorizing that has so constrained managerial analysis since the rise of the decision-making paradigm, and which KM may now—at last—be beginning to help us get beyond. It reminds us yet again of Simon's pivotal and paradoxical position; on the one hand a colossal contributor to the hegemony of the decision-making school, but on the other, one of a handful with the intellectual integrity and courage to point out from its very beginning that management would remain a mystery until we find ways of going beyond a rational decision making framework that cannot distinguish managers from appropriately programmed computers.

Concluding Remarks

Many will protest these books are totally unrepresentative of the state of KM, and they might point instead to the many IT-based conferences, books, journals and, especially, to the vast public and private sector IT spend. Others might point to the vigorous debates in human resource management and accounting—about which we should certainly know more than most of us do. Others might point to the academic work exemplified by the 2002 *Special Issue of Organization Science*, the *Strategic Management Journal*, the *Journal of Knowledge Management*, the *Journal of Management and Governance*, *Management Learning*, and so on. Yet others might argue that none of these approaches are capable of addressing the questions asked and that we should turn instead to quite differently grounded discussions within developmental psychology, parallel computation, neuro-biology, theories of social justice, chaos theory, thermodynamics, and so forth. This is one of the downsides of using the term knowledge which, by definition, applies to anything and everything.

It is tempting to contain these difficulties with a definitional sentence that begins with: 'what we mean by KM is ...' as if that could resolve the matter. It cannot, because KM is so intimately bound up with the notion of human agency, what Child called 'strategic choice'. We are only interested in KM because we live within a culturally contingent system of beliefs, one of which is that we have options, a measure of free will, a notion that infuses our world view at every level of analysis. We do not know how the world constrains our options. A wide range of managerial theorizing presumes we can know the world well enough to manage our choices 'rationally' and that the knowing being brought into the decision-making process is relatively unproblematic. But Simon told us that under such circumstances there is little need for managers or none for an administrative science.

As a theoretical activity KM stands on the empirically observed weaknesses of the rational model, on the one hand that our knowledge of the world is faulty, on the other that we cannot compute perfectly. Simon dithered between these two definitions of bounded rationality perhaps recognizing there may be no difference between them. Because there is no theory of such weakness, no axioms that correspond to those of economic rational man, KM is essentially inductive or 'grounded'. We try to generalize from observed phenomena about, for instance, tacit skills or group learning. But, in spite of the research into decision-making bias, these are not yet part of a comprehensive theory of the way we make choices, individually or collectively. If we had a theory of tacit knowing as, say, pre-conscious knowing, and were able to build a theory of the relationship between pre-conscious

and conscious knowing, then it should provide a model we could substitute for the economic rationality that Simon criticized. We would have a knowledge-based theory of choice that might lead us to a knowledge-based theory of why firms exist.

The first kind of KM seeks to extend conventional economic or organizational theorizing, invoking intangible assets, or the processes that produce it, without at the same time questioning its notions of choice. In this sense it may well deny its own problematic. Why is the manager's decision-making with explicit knowledge inadequate? If the answer is that others' choices seem to take into account more than they can say, we create a paradox. The proposition that their tacit knowledge can be made explicit enough to meet the pre-requisites of a rational decision-making model is tantamount to saying bounded rationality is a condition that, through rational management, can be rendered transient and correctable. Alternatively we might argue that in practice managers act as if bounded rationality is correctable by an act of interpretation or sense-making. In this case, if the process of surfacing others' tacit knowledge can be cut off at any point what theory of choice is now implied? What determines the interpretation and is this choice consistent with rationality?

Nor are we ever likely to come up with a comprehensive definition of knowledge. On the one hand there is the philosophy of science notion of policing knowledge, distinguishing between true and false knowledge. We know this ultimately drives us towards the tautologies of realism or to the relativism of the interpretive position. Contemporary philosophizing rejects the notion of false knowledge as a contradiction. Statements can be incorrect for a variety of reasons, but eventually our knowledge, beliefs, and practices converge. Knowledge just is. It might be considered self-referencing, forcing us back to an empirical approach, seeing knowledge as evidenced in our choices, as an aspect of the way we experience the world. KM becomes distinguishable from conventional theorizing when it stands on empirically shaped assumptions about the ways managers, teams, learning communities, firms, industries, markets, institutions, regions, nations and so on choose in spite of bounded rationality. It is about creating analytic frameworks in which to explore the options we believe relevant to theorizing about management and organizations under empirically observed conditions of imperfect knowledge. The possibility of a comprehensive theory of such imperfections is as remote as a comprehensive theory of knowledge.

KM begins, for instance, when researchers turn from technical questions about IT's capability to move abstractions around, to consider precisely and empirically what is being moved, how it gets created by or integrated into social practice, the observable consequences of its being moved. For many, KM is simply about transferring and appropriating economic value from existing knowledge. Others focus on the creating new knowledge, and yet others seek March's balance between exploration and exploitation. These processes only become theoretically interesting when practice shows they are not transparent in the sense of being 'perfect'. For instance, data is not 'natural' or unmediated, it requires a prior system of categories in which it is captured. These categories exclude as well as include, so by definition 'data' is partial and flawed. What are the implications of these category choices for the firm? A comprehensive theory of the firm would indicate the correct categories, while a knowledge-based theory would show how the firm derives from the choices actually made and empirically observed. Likewise intellectual property rights issues are important because, in practice, it is difficult to secure 'perfect' title to knowledge.

These three books are important because they eschew simplifying assumptions under which these kinds of KM problems are resolved by, for instance, turning non-traditional knowledge assets into 'regular' rivalrous ones. They also deny the maneuver of treating learning, the production of knowledge assets, as a process that can be managed in a rational manner. They keep pushing beyond these defenses of the mainstream to scout the more radical ground where our theories must be reconstructed after replacing the familiar rational choice assumptions with empirically grounded ideas about how we choose.

Overall the purpose of this essay is to help us see that while KM is not necessarily radical or 'critical', it can become so. It might then offer us a way to take up and advance Simon's legacy. We probably need a rich mix of disciplines, illustrated by Amin and Cohendet's book; a profound re-assessment of our research methodologies, illustrated in Patriotta's work; and an aggressive search for new philosophical underpinnings of the type exemplified in Tsoukas and Mylonopoulos's volume. There is a new frontier here.

Notes and References

1. A. Pickering, *The Mangle of Practice: Time, Agency, and Science*, University of Chicago Press, Chicago, IL, 1995.

J. C. Spender
Visiting Professor
Open University Business School
UK