*Journal of Management and Governance* **5:** 307–330, 2001.



# Models of Cognition, the Contextualisation of Knowledge and Organisational Theory

EDWARD LORENZ
*Centre d'Etudes de l'Emploi 29, promenade Michel Simon 93166 Noisy-le-Grand, France (E-mail: Lorenz@alum.mit.edu)*

**Abstract.** This paper examines the importance of cognitive foundations for theories of organisational behaviour. Three different conceptions of human cognition and reasoning are examined: the information processing, situated learning and cultural-historical perspectives. The paper shows how each conception of cognition leads to a different understanding of organisational routines and organisational problem-solving, as well as to the adoption of a different empirical methodology for observing organisational behaviour and for testing hypotheses about the nature of routines and problem-solving. The paper demonstrates that of the three approaches to human cognition, only the cultural-historical one gives rise to an understanding of organisational knowledge as embedded within a wider cultural and institutional setting.

## 1. Introduction

This paper examines the implications of different conceptions of human cognition for understanding processes of knowledge use and development within organisations.[1] It begins with an investigation of the two cognitive theories that in my view have had the most significant impact on the organisational behaviour literature: the information processing approach associated with Newell and Simon (1958, 1972), and the situated learning perspective, often associated with the work of Lave (1988) and Suchman (1987). The former was translated into the realm of organisational theory by Simon and March (1958) and Cyert and March (1962), and has been most fully developed in the contemporary research programme in computational, organisational theory (COT) associated with such authors as Carley (1995) and Masuch (1992). The latter has been extended into the realm of organisational theory in research analysing organisational learning in terms of "communities of practice" (Brown and Duguid, 1991; Lave and Wenger, 1991; Snyder, 1997; Wenger, 1998a, 1998b).

Following this exercise, I turn to the cultural-historical approach to cognition associated in the U.S. with the research activities of the Laboratory of Comparative Human Cognition at the University of California-San Diego. Although this perspective on human knowledge and cognition has not had a significant impact on the organisational behaviour literature, some of its implications for organisa-

tional issues, including processes of organisational learning, have been explored by Edwin Hutchins (1995).

One reason for examining cognitive foundations has to do with theoretical cumulativeness. In the organisational behaviour and strategy literatures, use is made of various concepts pertaining to knowledge use and development: "single-loop vs. double-loop learning," "mode 1 vs. mode 2 knowledge," "tacit knowledge" and "implicit knowledge," "mental models" and "cognitive frames." As Bart Nooteboom has observed, without an underlying theory of knowledge and cognition, the relation between these various concepts remains obscure and there is little room for scholars to build on each other's research results:

> It seems that inspired by experience from practice, different scholars more or less independently re-discover the same ideas, give them new names, and develop their own "grounded theory". If the different terms do not refer to the same concept, what is their difference and how are they related? If we do not find out, both replication and criticism of research are difficult and we will continue to proceed in a fragmented, haphazard, non-cumulative fashion (Nooteboom, 2000, p. 50).

In order to demonstrate the importance of cognitive foundations for organisational theory, I consider the implications of the three approaches to human cognition I have identified for how we understand "organisational routines" and "organisational problem-solving." As the analysis in the body of the text shows, the information processing approach gives rise to an understanding of routines as symbolic expressions stored within the minds of the organisation's members. It leads to the conclusion that problem-solving can be understood independently of the social context within which the organisation operates.

While the situated knowledge approach takes the opposite stance and argues that knowledge and learning can only be understood in relation to an external context, it appears to be firmly committed to the view that local context is determinant. As my discussion of the "communities of practice" literature shows, this gives rise to a conception of routines and problem-solving as necessarily tied to activity in particular times and places. This approach appears to close the door on a serious enquiry into the basis for more generalised forms of knowledge that stretch across different local contexts.

While the cultural-historical perspective on cognition shares the view of the situated learning approach that external context matters, it overcomes some of the limitations of this latter tradition. Firstly, as my discussion of Hutchins' (1995) work illustrates, it addresses some of the concerns of those wedded to an information processing conception of cognition by relating the coordinated behaviour of teams to operations on symbolic representations. Secondly, it provides a framework for seeing how the routine and problem-solving activities of organisational members are embedded in a wider historical and cultural context stretching across particular times and places.

## 2. Organisations as Information Processing Systems

The use of computational methods in the modelling of various theoretical schemes is often justified in terms of the intractability of systems of complex equations, or in terms of the opportunities it provides for performing experiments which cannot easily be carried out within the confines of the laboratory. In the case of computational organisational theory, however, the justification for using these methods has a deeper epistemological basis. Fundamentally, it stems from the belief that organisations, much like the human mind and the computer, are complex information processing systems. Correspondingly, it is argued that much can be learned about how organisations use and develop knowledge by studying the characteristics of computer programmes.

Newell and Simon elaborated the information processing conception of human cognition in the 1950s (Newell, Shaw and Simon, 1958). It holds that knowledge consists of representations or collections of abstract symbols that are stored in the mind, and that human reasoning and problem-solving behaviours may be understood in terms of operations performed with these symbols. Newell and Simon later developed a more formal statement of this view in the idea that the basic architecture of human cognition is a physical symbol system:

> A physical symbol system consists of a set of entities, called symbols, which are physical patterns that can occur as components of another type of entity called expression (or symbol structure) . . . Besides these structures, the system also contains a collection of processes that operate on expressions to produce other expressions: processes of creation, modification, reproduction, and destruction. . . . A Physical Symbol System has the necessary and sufficient means for general, intelligent action (Newell and Simon, 1976, p. 116).[2]

The physical symbol system hypothesis, of course, has increasingly been challenged within cognitive science. It can no longer be seen as holding the dominant sway over the profession that it did in the 1970s. Probably the most well-known, alternative cognitive architecture to be proposed is the connectionist one, in which knowledge is represented by a large network of parallel processes that store relations between features of inputs and associated outputs, rather than symbolic representations or descriptions.[3]

My concern here, though, is not especially with the question of whether the architecture of human thinking is symbolic in nature. Rather, I am concerned with another feature of the information processing approach to human cognition, namely, the idea that knowledge consists in representations that are literally *stored* within the mind, presumably in long-term or procedural memory. These representations arrive in the mind via "transducers" that transform external stimuli into internal representations. Problem-solving, for example, is typically understood in terms of search processes designed to find a match between external stimuli and the existing repertory of stored representations.[4]

The upshot of the stored representation view, as Hutchin's (1995, Ch. 9) has observed, is to restrict the unit of cognitive analysis to the individual mind and to view intelligence and human problem-solving as internal phenomena occurring in isolation from the external world of physical artifacts and social arrangements. As Newell et al. (1989, p. 107) put it, "Symbol systems are an interior milieu, protected from the external world, in which information processing in the service of the organism can proceed."[5]

This basic understanding of human reasoning and problem solving was extended to the realm of organisational decision-making by Simon and his colleagues at the Carnegie Institute of Technology in the 1950s and 1960s. Simon and March, in their seminal 1958 volume, *Organisations*, laid out a view of organisational decision-making based on the stored representation view of human cognition. In Chapter 6, focussing on the cognitive foundations of organisational decision-making, the basic unit of analysis is defined as a computing routine or a programme:

> Situations in which a relatively simple stimulus sets off an elaborate programme of activity without any apparent interval of search, problem-solving, or choice are not rare. They account for a very large part of the behaviour of all persons, and for almost all the behaviour of persons in relatively routine positions. Most behaviour, and particularly most behaviour in organisations, is governed by performance programmes (March and Simon, [1958] 1993, p. 163).

As the above citation shows, programmes are distinguished from the behaviour they govern. Simon and March clearly locate programmes within the minds of the organisation's members, while recognising that these internal representations may be written down with varying degrees of accuracy:

> Most programmes are stored in the minds of the employees who carry them out, or in the minds of their superiors, subordinates, or associates. . . . Programmes may be written down, more or less completely and more or less accurately (ibid., pp. 163–64).

The hierarchical nature of relations among the organisation's employees are paralleled by a hierarchical structure of programmes, in the sense that the programmes of superiors have as their main output the initiation or modification of the programmes of subordinates (ibid., p. 171). Specialisation of programmes among employees results in interdependencies among organisational sub-units. In relatively stable environments, the resulting burden of coordination can be incorporated directly into the programmes when they are established, thus reducing the need for on-going communication and feedback. In the case of unanticipated contingencies, coordination will require communication in order to adjust behaviour to deviations. While hierarchy is important in initiating and legitimising programmes, the communication involved in such feedback processes generally does not follow the line of command (ibid., pp. 180–184).

Consistent with the information processing view of individual cognition, problem-solving in the behavioural theory of the firm is analysed in terms of internal search processes designed to select an appropriate programme from the organisation's existing repertory. Organisational learning, on the other hand, is defined as the process of modifying the existing repertory or adding new programmes to it:

> If an organisation has a repertory of programmes, then it is adaptive in the short-run, insofar as it has procedures for selecting from this repertory a programme appropriate to each specific situation that arises. The processes used to select an appropriate programme are the "fulcrum" on which short-run adaptiveness rests. If, now, the organisation has processes for adding to its repertory of programmes or for modifying programmes in the repertory, these processes become still more basic fulcra for accomplishing longer-run adaptiveness. Short run adaptiveness corresponds to what we ordinarily call problem-solving, and long-run adaptiveness to learning (ibid., p. 192).

Since the basic unit of organisational analysis in March and Simon's approach is a cognitive one, namely, a programme or rule-like representation stored in the minds of the organisation's members, undertaking empirical work on organisations poses a problem of observation. This difficulty is fully recognised by the authors, and in the concluding postscript to the volume they identify the problem of observation as an important area for future research (ibid., pp. 234–235).

One way in which recent research within the information processing tradition has attempted to identify the programmes underlying routinised and problem-solving behaviours is through laboratory experiments. Perhaps the best-known work of this nature are the experiments carried out by Cohen and Bacdayan (1994), or the experiments directly inspired by their work carried out by Egidi (1996). In the appendix of a recent Santa Fe Institute working paper on organisational routines (Cohen et al., 1995), Egidi argues that the techniques of experimental psychology, including protocol analysis, allow for observation at the level of stored representations:

> If we suppose individuals are capable of forming an internal, mental representation of the situation, based on symbols and their manipulation, we can find experiments to confirm or refute our approach ... an obvious problem is to define to what extent a sequence of behaviors can be defined as "the same" ... most of the problems involved are solvable if we accept the idea of rule-based actions: if we can prove that individuals use the same set of condition-action rules during their activity (by protocol analysis, for example) then we can identify the routinised behaviors as "the same," even though we observe behavioral variation over time (Edigi, 1995).

Another, more widely applied, method of investigating the programmes governing organisational behaviour is the use of computational methods. The basic epistemological justification for this approach was laid out by Newell, Shaw and Simon

in their classic 1958 article, "Elements of a Theory of Human Problem Solving." This piece provides a clear statement of the epistemological stance regarding the adequacy of computer programmes for studying the properties of real-life problem-solving. The authors are careful to specify that the correlation they draw between computer programmes and human problem solving does not go all the way down. As they note, they "are not comparing computer structures with brains, nor electrical relays with synapses" (p. 8). Rather, the argument is that the computer and the human brain constitute two different physical supports or types of hardware for executing the same type of information processes:[6]

> Digital computers ... can, with appropriate programming, be induced to execute the same sequences of information processes that humans execute when they are solving problems. Hence, as we shall see, these programmes describe both human and machine problem solving at the level of information processes (Newell, Shaw and Simon, [1958] 1989, p. 8).

In a somewhat circular manner, the evidence assembled to prove this proposition was the output of a computer programme, the Logic Theorist (LT), which the authors had designed to solve simple problems in logic. Thus, their ability to programme a computer to solve certain sorts of problems in a way that mimed some of the known stylised facts concerning how humans solve such problems, was offered as proof of a theory of human cognitive processes inspired by the operation of a computer.

Simon explicitly referred in this research in his 1957 address with Newell to the Operations Research Association of America, where he argued, "We are now poised for a great advance that will bring the digital computer and the tools of mathematics and the behavioural sciences to bear on the very core of managerial activity – on the exercise of judgement and intuition; on the processes of making complex decisions" (Simon and Newell, [1957] 1982, p. 382).

The 1960s saw a first wave of computational models of organisational behaviour inspired by this research agenda, starting with those published in Cyert and March (1963), and culminating in the famous simulation of the "garbage can" model of ambiguity in organisational decision-making by Cohen, March and Olsen (1972). Despite the interest that this work created around the issues of bounded rationality and complexity in organisational decision making, it did not generate a sustained programme of research in computational organisational theory. In a recent comment on the sporadic advance of computational modelling, Cohen (1998, p. ix) suggests that part of the explanation may lie in the poor match that existed between the control processes of traditional computing languages (typically FORTRAN), involving iterative loops and conditional branching, and actually occurring organisational coordination mechanisms. This was arguably no where more evident than in his own "garbage can" model, which as a number of critics have observed is curiously devoid of organisational structure (Padgett, 1980; Carley, 1986).

The late 1980s marked the beginning of a new wave of computational models in the field of organisational behaviour. These recent modelling efforts display in various ways the influence of prior developments in the field of artificial intelligence. Firstly, as with the symbolic, artificial intelligence community more generally, the information processing or physical symbol system hypothesis is regularly invoked to justify the use of computational methods. Preitula, Carley and Les Gasser (1998, p. xiii) provide a clear statement of this position in the introduction to their recently edited volume on computer simulation, where they argue that "Computational organisational theory (COT) is the study of organisations as computational entities. COT researchers view organisations as inherently computational, because they are complex information processing entities."[7]

Secondly, the current wave of modelling exercises has employed recent techniques and concepts developed in artificial intelligence. These include the use of genetic algorithms to explore adaptive learning processes, work in distributed artificial intelligence, and object orientation in programming. The latter two of these, as Cohen (1998, p. x) observes, respond to some of the limitations noted above of traditional programming languages. These developments have arguably given new support to the idea that a realistic account of knowledge use and decision-making processes within organisations can be provided by studying the properties of computer programmes.

> Objects provide a highly natural way of implementing model agents that have specialised capabilities and subtle, implicit, networks of interaction. This is a far more congenial framework within which to express intuitions about organisational processes. ... A second development is the growing interest in organisational questions shown by researchers in the field of distributed artificial intelligence, who have recognized that distributed computers and programmes must deal with many issues that are profoundly similar to those facing human organisations (Cohen, 1998, p. x).

Simon and March's "behavioural theory of the firm" translated the information processing view of human cognition onto the organisational level. The underlying cognitive framework determines the basic unit of organisational analysis, and it guides the choice of empirical methodology used to test the theory.

As I observed above, one reason for the declining popularity of the information-processing view of cognition is the limited success of artificial-intelligence techniques based on it, in accomplishing such apparently simple learning tasks as face recognition. Another reason is the accumulation of neuro-biological evidence, indicating that the sensorimotor system is a locus of cognitive activity rather than just a mediator between the external environment and the mind.[8] This, in turn, has done much to undermine Simon's position that we can safely ignore the specifics of the "hardware" in studying the information processes that are presumed to govern human reasoning and behaviour.[9] A third reason is the variety of ethnographic evidence showing how basic cognitive aptitudes, such as the ability to carry out simple mathematical operations, vary depending on the environmental context in

which the individual performs. This kind of evidence has been used to support a "situated" perspective on human cognition that, contrary to the stored representation view, argues that knowledge and learning can only be understood in relation to a local actional context. In the following section, I turn to the situated action perspective and the "communities of practice" approach to organisational behaviour and learning, which it has inspired.

## 3. Situated Action and the Communities of Practice Approach

The notion of a "community of practice" grew out of the work on the situated nature of learning by John Seely Brown, Paul Duguid, Jean Lave, Lucy Suchman and their associates at the Palo Alto Institute for Research on Learning in the 1980s. In terms of its intellectual roots, it is probably fair to say that sociological theories of situated practice and action were more important influences on their work than psychological theories. For example, the methodological premises of Chicago school symbolic interactionism, associated with Herbert Blumer, have much in common with the view of Brown and Suchman, that knowledge is necessarily tied to an actional context. Thus, Blumer (1969, p. 2) argued that the meanings one attaches to the behaviour of others arises out of social interaction, and that these meanings are modified through an interpretative process in the course of dealing with others. The reliance on ethnographic methods similarly follows in the tradition of Chicago school qualitative sociology, notably the work of Robert Park and W.I. Thomas. A focus on the everyday or lay methods used by individuals to make sense of what others do and say as a tool for exploring the nature of knowledge and representation, especially present in Suchman (1987), draws inspiration from Garfinkel's (1967) ethnomethodology.

While the dominant, intellectual roots of the situated learning perspective can be found in sociology, there was a clear concern to draw out the implications for mainstream cognitive science and for theories of knowledge and learning inspired by it. The challenge they posed to the mainstream approach was to contest the view that knowledge consists of representations that are literally stored within the mind, presumably in long-term or procedural memory, and that intelligence and problem-solving capabilities may be understood in terms of formal operations on those representations. The alternative view they elaborated held that knowledge and problem-solving abilities are necessarily situated relative to a local context that includes not only people's interrelations with one another, but also artifacts. The way the artifactual context is experienced may differ from one individual to another. For this reason, the structural framework in which people act is necessarily constructed and malleable (Lave, Murtaugh and de la Rocha, 1984, pp. 71–72).

In developing this view, researchers at the Institute for Research on Learning mostly by-passed the question of the internal architecture of human cognition. In particular, they did not seek to challenge the mainstream view by mustering the psychological or neurobiological evidence that might support an alterna-

tive description of internal cognitive architecture, in the sense that Newell and Simon used the term.[10] Rather, they challenged the mainstream perspective more indirectly by providing field evidence of how cognitive processes and problem-solving capacities depend on a subject's interaction with the external environment. For example, Lave, Murtaugh, and de la Rocha (1984)[11] find that arithmetic procedures used on the job or in one's daily activities, such as in shopping, are structured and aided by features of the physical setting in which they occur. These skills transfer poorly or not at all onto different contextual settings, such as the classroom. Suchman's (1987) ethnographic account of the difficulties novice users faced in their interactions with a computer-based help system attached to a complex photocopier shows how language and words acquire meaning or are "indexicalised" by the shared context in which they are used. In this sense, one can say that the shared environment contributes to the intelligibility of conversation.[12]

Suchman's (1998) recent study of problem-solving by the personnel of the operations room of an airline at a mid-size airport in the United States provides another illustration of the contribution of local context to problem-solving activity. As she observes, the episode she focuses on can be viewed as an instance of routine trouble and problem-solving. It concerns the operations room of an airline in a mid-size airport in the United States. The personnel of the operations room have responsibility for coordinating the ground activities involved in servicing in-coming and departing aircraft. The problem situation she analyses concerns the faulty operation of a set of stairs that are wheeled onto a gate to service an arriving airplane. The problem is eventually solved by the airline's ground maintenance personnel, and Suchman's account focuses on the efforts of the operations personnel to coordinate that process, which included making requests to borrow the stairs of another airline operating at the same airport.

The operations personnel work according to a division of labour with different employees responsible for such tasks as monitoring booking and connections, assuring communication with the pilots of incoming and departing aircraft, and for advising the ramp crew of the arrival of incoming flights. What Suchman's account demonstrates is that this division of labour is embedded in a particular set of spatio-orientational arrangements, and that coordination of activities among the personnel depends on spatially-oriented shifts in gaze and changes in body position. Such oriented movements and shifts of gaze can be used by workers to tacitly indicate a partial or total reorientation of their attention to joint-problem-solving activities. The analysis shows how spatio-orientational arrangements are critical to acheiving the "shared focus of attention" (p. 57), which problem-solving under conditions of specialisation and the division of labour requires.

Suchman's account displays many of the hallmarks of the situated-action approach to explaining human problem-solving capabilities. Such capabilities emerge out of people's actual practice in particular, artifactual settings and the knowledge remains both highly-contextualised and tacit in nature. It is precisely this highly contextualised and tacit nature of their knowledge, though, that makes

the insights of situated-learning approach difficult to apply to the field of organisational behaviour. One might well ask: What general conclusions can be drawn from this work for the relation between an organisation's structure and its problem-solving capabilities? The notion of "communities of practice" was developed, at least in part, in an effort to link organisational structure to organisational knowledge and problem-solving.

The basic idea is that communities of practice consist of people bound together by informal relations who share a common practice. Around this shared practice, they develop a common language and shared understandings of the environmental context in which they work, including the meanings they attach to artifactual arrangements. These shared understandings remain tacit or implicit for the most part, and communities of practice are characterised by a low degree of codification of members' knowledge (Brown and Duguid, 1991, pp. 41–43; Wenger, 1998b).

Despite their informality, communities of practice have boundaries and behaviour within them that is routinised. By structurally embedding knowledge in this manner, it becomes possible to use the notion of communities of practice to address some of the concerns of organisational behaviour theorists. Researchers such as Brown and Duguid (1991), and Wenger (1998b) and Snyder (1997) instruct us that firms consist in collections of possibly overlapping communities of practice, and that these communities are central to both the transmission of knowledge within organisations and to the organisation's problem-solving capabilities. As Wenger (1998b) recently put it, "they [communities of practice] are a company's most versatile and dynamic knowledge resource and form the basis of an organisation's ability to know and learn."

How do communities of practice serve these functions? Firstly, they provide a mechanism for the transmission of tacit knowledge and the reproduction of routinised behaviours. The key concept here is that of "legitimate peripheral participation," introduced by Lave and Wenger (1991).[13] The basic idea is that people acquire the knowledge they need to achieve a competent performance by becoming "insiders" or legitimate members of the relevant community. Initially, as newcomers, they are at the periphery of the community. They observe the behaviour of old-timers and they acquire the language and worldview of the community by participating in its discussions. Over time, they become knowing and fully skilled members of the community, capable of participating in the process of passing on the community's knowledge to others. Thus, learning is eminently a process of enculturation. As Brown and Duguid (1991, p. 48) put it, "learners are acquiring not explicit, formal 'expert knowledge,' but the embodied ability to behave as community members."

Secondly, the notion of communities of practice provides an account of problem-solving capabilities, which sharply contrasts with the stored-representation view of problem-solving as a search process designed to select an appropriate programme from an existing repertory. The basic idea is that confronted with an organisational problem, say an equipment malfunction,

members of a community engage in narration or story-telling. The stories narrate the members' past experience with similar problems and they are designed to draw on the community's accumulated wisdom in an effort to arrive at a shared diagnosis of the problem. As Brown and Duguid (1991, p. 44) have observed, this form of problem-solving is quite different from following the branches of a decision tree in an explicit programme or solution algorithm: "They (stories) function, rather like the common law, as a usefully underconstrained means to interpret each new situation in the light of accumulated wisdom and constantly changing circumstances."

The idea that communities of practice are self-organising bodies whose members share practices tied to particular artifactual settings makes it a powerful concept for understanding how organisational groups use tacit knowledge in their daily problem-solving activities. The epistemological implications of the approach directly follow the underlying situated knowledge perspective on human cognition. If we want to improve our understanding of organisational knowledge use and problem-solving, we are well advised to adopt an ethnographic approach and to focus on observing the informal forms of association among members that cut across the firm's formal administrative job and occupational descriptions.

The emphasis on the local contextualisation of knowledge in the research on communities of practice also accounts for certain limitations of the approach, especially as it has been developed in recent literature directed at practitioners.[14] Firstly, while this literature informs us that communities of practice may cross over company boundaries, little attention is given to the formal, professional craft or artisan associations that often cater to the interests of such wide-flung communities. Formal organisation is often used by such occupational or craft communities to regulate apprenticeship and to provide a forum for the informal exchange of ideas and knowledge that can lead to the diffusion of organisational practices over wide geographic spaces.[15] More generally, and closely related to this, there is little effort in the communities of practice literature to consider how local, contextualised forms of knowledge are articulated and interact with more widely held forms of knowledge that stretch across particular times and places.[16] For example, little attention is give to the forms of shared training and education outside the workplace that may underpin norms and conventions that bear directly on expectations and behaviours at workplaces spread over wide geographic areas.

To do justice to the communities of practice research agenda, though, the question of the relation of the global to the local, or of the general to the particular, is not entirely ignored. Some effort is made to carve out a middle-ground, as in the following quote from Wenger's 1998 volume, *Communities of Practice*:

> What transpires is that knowing is defined only in the context of specific practices, where it arises out of the combination of a regime of competence and an experience of meaning. Our knowing – even of the most unexceptional kind – is always too big, to rich, too ancient, and too connected for us to be the source of it individually. At the same time, our knowing – even of the most elevated

> kind – is too engaged, too precise, too tailored, too active, and too experiential for it to be just of a generic size (Wenger, 1998a).[17]

One can see in this quote a certain tension in how the relation between the general and specific are handled. On the one hand, there is a desire to recognise the reality of generally held representations and rules, but there is simultaneously a desire to argue that such things only have meaning in relation to peoples' local contextualised experience. The local level, then, is what ultimately determines the way ideas and principles are formed and take shape.[18] This, in turn, would appear to discourage any interest in undertaking a serious inquiry into the mechanisms that might serve to transmit and diffuse knowledge beyond the confines of particular places and times. At the end of the day, the stance elaborated in the communities of practice perspective can only lead one to doubt the relevance of such notions as societal norms and conventions for an analysis of organisational behaviour and problem-solving.

In what follows, I will argue that the work stemming from the activities of the Laboratory on Comparative Human Cognition in the 1980s takes an important step towards overcoming these limitations by examining the way artifactually-mediated cognition, including language, tie local interactions to wider social, institutional processes.

## 4. The Cultural-Historical Approach to Distributed Cognition

Michael Cole, Yrjö Engeström and others associated with the work of the UC-San Diego Laboratory on Comparative Human Cognition (LCHC), such as Edwin Hutchins, share the view of researchers at the Institute for Research on Learning that knowledge is necessarily situated relative to a context. Indeed, it is impossible to draw a neat division between these two communities of scholars. Jean Lave, for example, who developed the concept of "communites of practice" with Etienne Wenger was closely associated with the activities of the LCHC in the early 1980s.

What distinguishes the work of researchers associated with the LCHC, though, is their emphasis on the cultural and historical determinants of cognitive processes. This emphasis leads them to tie local context to wider social and institutional arrangements in a way that is absent in work inspired by sociological theories of situated practice.[19] To a large extent, this can be accounted for by the influence of sociohistorical psychology or "activity" theory, associated with Vygotsky and Leont'ev. Consistent with these intellectual roots, much of the work of the LCHC has focussed on problems of child development, such as the development of literacy or basic perceptual skills. The closest thing to an application to the field of organisational behaviour comes in the work of Edwin Hutchins (1995). For this reason, after developing a few of the core concepts, I will turn to Hutchins' recent study of problem solving by navigation teams.

Perhaps the core idea in the sociohistorical approach to psychology is that human cognitive processes are mediated by tools and artifacts.[20] Tools and artifacts

are taken to refer not only to physical objects, such as the pencil and paper I may use to make a calculation or the computer I may use to produce written text, but also to language and external symbolic representations in the form of verbal texts and the like. As an example of the latter kind of tool, think of the engineering process books that serve as aids in product assembly. Such tools not only transform the nature of the cognitive processes needed to perform various tasks or computations, but also serve to connect the individual to knowledge held more widely in society. As Leont'ev put it in his summary of Vygotsky's work:

> The tool mediates activity and thus connects humans not only with the world of objects, but also with other people. Because of this, human activity *assimilates the experience of humankind*. This means that human's mental processes (their "higher psychological functions") acquire a structure necessarily tied to sociohistorically-formed means and methods transmitted to them by others in the process of cooperative labor and social interaction (Leont'ev [1981] as cited in Werstch, [1987] 1997, pp. 226–227).

Another key concept is that of the "zone of proximal development." This was developed to describe the way children benefit from their interaction with adults, or more experienced children in the development of their cognitive and problem-solving capacities. Initially, the child's participation in task accomplishment may be quite minimal, amounting to little more than recognition that a task is at hand. As an LCHC paper observes, by participating in a task that is understood by the adult "a child can 'accomplish' a task before the child understands what he or she is doing." Over time, as the child comes to understand the reasons for the various actions involved in accomplishing the task, there occurs a process of internalisation of the regulative speech and external symbolic expressions that mediate task performance for the adult. This implies that, "the child's understanding of the task and of the associated, complex-regulative speech of the adult is a consequence, rather than a prerequisite, of going through the task" (LCHC, 1983, p. 336). Vygotsky (1979, p. 163) summed up this idea in his so-called "general genetic law of cultural development," stating that the intrapsychological plane, including concept formation, appears first on the interpsychological plane in the form of social relations.[21]

These two ideas, reinterpreted in the contemporary language of cognitive anthropology and psychology, are developed in a masterful way by Hutchins (1995) to give an account of the cognitive properties of the work of navigators on a US Navy vessel, responsible for determining, charting and logging the ship's position at various time intervals. Navigation work, at least in large vessels such as that studied by Hutchins, is distributed across the members of a team. In the case studied by Hutchins, the process of determining the ship's position necessarily begins with specialised operators observing visual bearings of landmarks with a special telescopic-sightings device called an alidade. These visual bearings are then progressively translated into a series of angular and digital representations by different personnel working with the aid of various physical artifacts in accordance

with a well specified algorithm. Ultimately, the angle measured between the ship and the landmark in the world is reproduced as a representation in the space of the latitude and longitude grid of the ship's charts.

In keeping with Vygotsky's notion of tool-mediated cognition, Hutchins analyses this process in terms of the propagation of representational states across a series of representational media or artifacts. The mediating artifacts include both physical objects and tools (e.g. the alidade and the chart), and the cognitive structure within the minds of the members of the navigation team. The representations of the ship's position take different forms in the different media: the angular representation in the alidade, the digital representation in the bearing record log, the memory of that representation in the mind of the navigation bearing recorder, etc. Representational states are propagated from one media to another by bringing the states of media into coordination with one another. As Hutchins observes, the tools and other media that are brought into coordination permit difficult tasks to be transformed into ones that can be done with the manipulation of simple physical systems, or possibly by the mental simulation of the manipulation of simple physical systems (ibid., pp. 171–72). This implies that the cognitive properties (e.g. capacity for error-detection and problem-solving) of the system made of team members in interaction with various artifactual media will differ from the cognitive properties of the individual members.

Performance in this cognitive system depends not only upon the use of physical tools, but also on the way task realisation is mediated by cognitive artifacts. For the apprentice navigator, this might be an external cognitive artifact in the form of a written procedure used as an instructional aid in the context of a training session with a more experienced navigator. However, with some experience, the novice will memorise the written procedure and it will exist as an explicit representation within the media of his internal memory. In this case, the mediating cognitive artifact is "internal rather than external" (ibid., p. 303). It can be seen as the cognitive residue of the novice's interaction with the external artifact and the more experienced team member. With even more experience the procedure may become tacit or implicit, embedded within the media of the individual's sensorimotor system. At this point, as is common in expert performance, the skilled navigator may find it difficult to account for his own task performance (ibid., pp. 310–311).[22]

Much in keeping with Vygotsky's concept of the "zone of proximal development," Hutchins' account of navigation apprenticeship shows how internal cognitive structure is shaped by the individual's insertion into a set of social relations.[23] Hutchins observes (ibid., p. 282) that a good deal of the cognitive structure that the novice will have to learn in order to perform the various steps in fixing the ship's position is present in the organisational relations among the members of the navigation team. Thus, in coming to understand the way the various members of the team depend on each other, the novice is simultaneously learning about the nature of the computation being performed and the ways the various parts of it depend on one another.

Perhaps the most well-know part of Hutchins' study is his striking account (Ch. 8) of the adaptive response of the Navy vessel's navigation team to an emergency situation occasioned by the failure of the ship's propulsion system while undertaking entry to the San Diego harbour. The loss of propulsion led to various electrical failures resulting in a loss of power to the ship's gyrocompass and the navigation team's dependence on visual readings from a magnetic compass which is subject to error. The team had no prior experience with coordinating the fix cycle in such circumstances. Hutchins' reconstruction of this emergency, based on video tape recordings, shows how through a trial and error process of mutual adjustment, the navigation team managed to put in place a new division of labour and make the necessary computation to determine the ship's position. As he observes, none of the members appeared to reflect on the process as a whole. Rather, team members put constraints on each other's behaviour through presenting each other with partial computations to which they mutually adapted.

The analysis brings out in an exemplary manner the idea that the sense an individual attaches to what others do and say within a distributed task is negotiated through on-going interpretative acts. This, of course, is one of the key themes in situated action theory, and this account of the navigation emergency might be taken as support for the idea that knowledge and social order are locally constructed, independently of wider institutional arrangements and norms. This, however, fails to do justice to the historical and cultural dimensions of Hutchins' account.

In the concluding chapter to his book, Hutchins makes clear that the relatively fleeting processes of adaptation and interpretation described in the navigation emergency occur within the context of the longer-term, cultural transmission of knowledge that is "crystallised and saved in the physical and conceptual tools of the trade and in the social organisation of work" (ibid., p. 374).[24] By social organisation, he is thinking of the way occupational roles and career structures, such as those described for the navigation team, are the products of an organisation with its own history and traditions, in this case the U.S. Navy.[25] As I observed above, this social organisation conditions the process of acquiring skills and knowledge. By conceptual tools, Hutchins is thinking of culturally rooted cognitive artifacts, both internal and external, which influence the way the problem is represented and the nature of the algorithms used in its solution. This aspect of his analysis comes out most clearly in the comparisons he makes between the representational assumptions of western and micronesian navigational practice (Ch. 2). The comparison brings out in a striking manner how even such basic premises embodied in our problem-solving representations, such as viewing the vessel in motion relative to externally fixed landmarks, rather than seeing the landmarks in motion relative to a vessel whose position is unchanging, are products of a particular culture and history.

### 4.1. IMPLICATIONS FOR ORGANISATIONAL BEHAVIOUR

The historical-cultural approach to cognition implies that the routinised and problem-solving behaviours of organisational members are emergent features of their interaction in carrying out distributed tasks with the help of external and internal mediating devices. By external mediating devices, I am thinking not only of tools, equipment and plant layout, but also of verbal texts such as standard operating procedures and instruction manuals that provide some sort of verbal description of the performance of tasks and skills. By internal cognitive artifacts, I am referring to the internalised representations of these texts in the minds of the organisation's members.

A key point developed in the historical-cultural account is that these mediating artifacts and devices are social products. They are produced by society's organisations and institutions with their distinctive histories and cultures. For this reason, their use by an individual in a particular instance of task performance serves to connect that individual to society more generally, and provides a basis for shared understanding and knowledge that extend beyond that individual's interactions at particular locations and times. This, however, is not something that can be explained in terms of the artifacts alone, but rather by the experience of their use in solving particular types of organisational problems.

Consider a mediating artifact such as the engineering process books used as an aide in product assembly. The argument is not that such mediating devices serve to transfer knowledge directly across different contexts. The cultural-historical approach makes a clear distinction between the peformance descriptions of a task given in a mediating artifact and the meanings the task performer attaches to those verbal descriptions. Meanings are not given by the language of the text but rather come about through the task performer's experience in carrying out the task, which involves coordinating the mediating artifact with the task environment. For this reason, language and verbal texts, rather than being thought of as instances of codified knowledge, are better thought of as coordinating tools used in task performance. Of course, with considerable experience the task performer will be able to imagine the meanings of the action descriptions given in the mediating artifact in various real word contexts and to form a mental model of the task at hand. This again raises the basic idea of Vygotsky that symbolic expressions and tokens initially exist in the external environment and only subsequently in the mind.

The idea that language and verbal text serve as tools for bringing structure within the mind of the task performer into coordination with the task world also implies a distinction between the descriptions of the performance of a skill given in the artifact and the actions taken on the world. Mediating artifacts should not be seen as devices that serve to directly transfer organisational routines by virtue of the way they directly control behaviour. Rather, as Hutchins ([1986] 1997, p. 338) has argued, such mediating devices typically set up constraints that serve to control behaviour indirectly by allowing the actor to evaluate his or her behaviour and to judge whether or not it is appropriate. Hutchins' example of a check list used in

carrying out an ordered sequence of actions in accomplishing a task brings out this point nicely. The check list allows the task performer to determine whether it is appropriate to move on to the next step, while typically leaving him or her considerable discretion in determining how each step should be carried out.

The use of common artifacts by individuals working in different physical settings to solve similar production problems provides a basis for the emergence of shared understandings and knowledge. However, this does not imply that behaviours will be exactly reproduced across contexts. By way of example, consider the knowledge-transfer problem confronting any large multi-plant enterprise establishing a production site at a new location. The transfer process will typically involve the use of such mediating artifacts as training manuals and process books that are used by the personnel originating from the parent firm as aids in transferring their know-how and methods to the employees of the new affiliate. As detailed empirical comparisons of such multi-plant enterprises have shown, there is invariably much that is idiosyncratic and specific to the operations of the various production sites. Nonetheless, the experience of using common artifacts to solve production problems which are common to the enterprise as a whole (e.g. assembling the same product), leads to the emergence of shared knowledge and understanding that make the behaviour of the members of one plant of the enterprise largely intelligible to the members of another plant.

Similarities in organisational routines and problem-solving behaviours are not limited to the different sites of multi-plant enterprises. There is vast literature pointing to national specificity in methods of work organisation and management.[26] For example, it is common to refer to distinctively Japanese methods of production (Aoki, 1988; Dore, 1973; Abegglen and Stalk, 1985). This does not imply the absence of significant diversity across enterprises in Japan, even within the same industrial sector. Rather, it refers to common features of the way work is organised and the way production and quality problems are solved, which appear particularly salient when Japanese firms are compared with their U.S. or European counterparts.

To some extent, such national specificity can be attributed to differences in the way national governments regulate labour markets, including occupational boundaries and the certification of skills and competences. Yet, as detailed, comparative case-study work makes apparent, there are many similarities falling under the heading of informal or "noncanonical" practices that cannot be linked in an obvious way to differences in a nation's legal framework and the actions of its regulatory agencies. The forms of communication and information exchange displayed by team members in joint-problem-solving activity provide a good example.[27] The fact that such coordinated behaviour by experienced team members is unmediated by written documents or tools often leads to the conclusion that the performance is a pure result of highly localised and context-specific process of sense-making and interpretation. Without denying the importance of unplanned processes of local adaptation and interpretation, it nonetheless remains the case that verbal texts and

documents may have played a role in the forms of initial training and education received by the team members. The role played by such external cognitive artifacts may no longer be apparent, due to their subsequent internalisation by the individuals in controlling their own behaviour. Shared knowledge and common ways of working across enterprises may reflect, in part, the influence of artifactually-mediated forms of initial educational and training that are common to various occupational groups and categories.

Further, as Hutchins' ([1986] 1997, pp. 351–352) account of mediation and internalisation suggests, this suggests a plausible explanation for the difficulty expert performers often display in explaining why they act in the ways they do. The reason is not that there never existed an explicit, internal symbolic representation of the performance of the task in the expert's mind, as may be the case for motor skills. Rather, the reason might be that the representation was closely tied to the mediating artifacts that were used as tools in the initial learning of the task, and that the memory of these mediating devices has atrophied. The performance is now automatic and unmediated, but the only way to account for it is to reconstruct through memory the experience of using the mediating devices, something which is now distant in time.

## 5. Conclusion

The introduction to this paper raised the question: Why is it important for organisational theory to examine cognitive foundations? The body of the text has provided some elements of a response to this question. It has contrasted the implications of the cultural-historical approach to knowledge and cognition for the study of organisational routines and problem-solving with the implications of the information-processing and situated-practice approaches. The information-processing perspective, as exemplified in computational-organisational theory, treats routinised behaviours as governed by programmes or symbolic expressions stored in the mind. The selective pressures accounting for routinised behaviours are internal to the mind, and much is explained in terms of reducing cognitive load. Environmental context is reduced to a stimulus that sets off these cognitive processes. Problem-solving in the information-processing account amounts to searching for an appropriate programme from the existing repertory, given an environmental stimulus.

In situated-action theory, as it is developed in the communities of practice literature, routinised behaviour emerges through a group's experience of a shared practice in a local context. Selective pressures are located outside the mind in the local environment composed of other actors and artifacts, which are essential "affordances" in the development of the shared knowledge and understandings that allow the members of a tight-knit community to communicate and make sense of each other's behaviour. Since knowledge is largely tacit, narration and story-telling are useful problem-solving tools. They help the members of the community to

articulate what they know implicitly, and they serve to transmit this knowledge among members within the community.

In the cultural-historical approach to cognition, reducing cognitive load and using the local environment to make sense of the words and actions of others can play a role in accounting for routinised behaviours. Selective pressures, however, may also act across time and space accounting for similar routines and behaviours in different contexts. One of the institutionalised mechanisms through which this may occur is a regulated system of apprenticeship, which serves to produce common language, skills and ways of doing things across firms. This sort of process was evident, for example, in Hutchins' account of the way the Navy's "culture," consisting of common occupational roles and training, generates similar behaviours across different vessels and local contexts (Hutchins, 1995, pp. 25–26).

Language and external-symbol tokens, such as instruction manuals or engineering process books, can also serve as tools in the emergence of shared knowledge and behaviours that stretch across particular contexts or localities. Thus, as the vast literature on the diffusion of organisational practices makes clear, codified descriptions and verbal characterisations of the performance of tasks are important mediating devices in the transmission of organisational knowledge and the reproduction of routines and behaviours. The fact that such devices may be no longer apparent in the case of a skilled performance does not imply that they never played a role. The noncanonical practices emphasised in the communities of practice literature may be grounded to some extent in the initial process of training and apprenticeship that relied on external mediating devices in order to develop the, "ability to carry out socially formulated and goal directed actions" (Wertsch, 1979, p. 32). This raises Vygotsky's (1979, pp. 157–163) view that symbolic expressions initially exist in the external world of social relations and only subsequently in the mind through processes of internalisation. This reverses the relation posited by the physical symbol system hypothesis between understanding and behaviour on the one hand, and internal symbol processing on the other.

The importance of socially-constructed mediating devices in the development of knowledge and the capacity for a competent performance in no way precludes that localised processes of adaptation and interpretation may contribute to the emergence of novel solutions and behaviours. A strong implication of the historical-cultural approach, though, is that such unplanned process of learning and adaptation occur within a broader context, shaped by society's more general norms and institutions.

## Notes

[1] Preliminary versions of this paper were presented at the 'Workshop on Cognition and Evolution in the Theory of the Firm,' Jena, Germany, September 25–27, 2000, the DRUID Winter Conference, Klarskovgaard, Denmark, January 18–20, 2001, and the DRUID Summer Conference, Aalborg, Denmark, 12-15 June 2001. I would like to thank the participants for their useful comments. This

paper has also benefited from the comments of Nicolai Foss and from numerous discussions with Christian Bessy about the work of Herbert Simon and Edwin Hutchins.

[2] Newell et al. (1989, p. 103) caution that the architecture should not be confused with a representation of the external world. The architecture supports such a representation but does not itself provide it.

[3] One of the appeals of the connectionist architecture is that machine learning methods inspired by it, notably neural nets using back propagation, have shown promise in accomplishing such apparently simple learning tasks as face recognition that have proven impossible for conventional AI (Gallant, 1993, Ch. 1).

[4] See Clancey (1997, pp. 46–75) for an extended discussion of this feature of the conventional approach.

[5] This is closely linked to the idea that the sensorimotor system serves as mediator between the world and the cognitive system rather than as an actual locus of cognitive activity. As Newell et al. (1989, p. 117) observed, "Perception and motor behaviour are assumed to take place in additional processing systems off stage. Input arrives in working (memory), which thus acts as a buffer between the unpredictable stream of environmental events and the cognitive system."

[6] For a discussion of these points in the context of a more general critique of the computational approach to understanding organisational routines, see Lazaric and Mangolte (1999).

[7] Their position is by no means idiosyncratic. See, for example, Blanning (1992) in his AI model of organisational decision making: "This hypothesis [human organisations are physical symbol systems] derives from the notion that organisational structure and knowledge, both implicit and explicit, impose a degree of formality in information processing and decision making that allows organisations, profitably, to be viewed as formal systems whose internal workings can be described in symbolic form." Also see Masuch and Lapotin (1989, p. 44), "This hypothesis [the physical symbol system hypothesis] – corroborated by the broader experience of AI models of decision making . . . – serves as the point of departure for the design of a new, generic model of organisational decision making."

[8] For a discussion of some of the evidence in the context of a more general discussion of the debate around codification of knowledge in the science and policy literature, see Nightingale (2001). Also see Clancey (1997, pp. 76–97 and 147–161).

[9] See, notably, Simon, 1996, pp. 13–24.

[10] See, however, Clancey's (1997) work in the area of robotics and artificial intelligence. He seeks to provide empirical evidence of a close coupling between perception, inference and action, rather than treating these three processes as occurring in separate and independent times as in the physical symbol system hypothesis.

[11] Also see, Lave ([1981] 1997).

[12] See Brown, Collins and Duguid (1989) for an overview of this and related literature.

[13] The idea of legitimate peripheral participation derives inspiration from Vygotsky's (1978) notion of the 'zone of proximal development.' Vygotsky developed this notion to characterise the way child development proceeds through the child's participation in activities slightly beyond his or her competence, with the assistance of adults or more competent children. See Laboratory of Comparative Human Cognition (1983, p. 338) and Rogoff (1990, pp. 14–15).

[14] See, for example, Wenger (1998b), Snyder (1997), and McDermott (1999).

[15] The emphasis on spontaneous and possibly unrecognised communities of practice in this recent literature contrasts with the examples of communities of practice in Lave and Wenger (1991), including tailors, butchers and quartermasters, that correspond to craft or occupational communities with well-defined rules for entry to the trade, based on regulated forms of apprenticeship.

[16] There is, of course, nothing new in this criticism. It is the standard one that has been levied in one way or another against symbolic interactionism, ethnomethodology and other currents of thought in sociology that ground their analysis in the situated nature of practice. See Fisher and Strauss (1978, pp. 485–486) for this traditional criticism raised against symbolic interactionism.

[17] In a similar vein, see Lave and Wenger, 1991, pp. 33–34: "Generality is often associated with abstract representations, with decontexualisation. But abstract representations are meaningless unless that can be made specific to the situation at hand. Moreover, the formation or acquisition of an abstract principle is itself a specific event in specific circumstances. Knowing a general rule by itself in no way assures that any generality it may carry is enabled in the specific circumstances in which it is relevant."

[18] For an uncompromising statement of the dominance of the local over the global, see Leigh Star (1998, p. 312): "We cannot find a rational, reliable resting place in the old way by falling back on ideas of routines, standards, and universal reliability – we are in the process of discovering that it's local all the way down, including the work of transporting it across locales."

[19] For a critique of 'context-specific approaches' and a discussion of how language can serve to carry over experience from one context to the next, see: Laboratory of Comparative Human Cognition (1983, pp. 329–332 and 340–342). For a similar emphasis on the role of language in collective knowing and memory, see Wertsch ([1987] 1997).

[20] For a discussion of the notion of tool mediated cognition in the work of Vygotsky, see Cole and Engeström (1993, pp. 4–10) and Wertsch (1979, pp. 23–26).

[21] Wertsch (1979, p. 32) provides a nice synthesis of Vygotsky's core ideas in observing that, "the notion of internalization is concerned with the ontogenesis of the ability to carry out socially-formulated, goal-directed actions with the help of mediating devices."

[22] Also see Hutchins ([1986] 1997) for a discussion of mediation and skill automatisation.

[23] See, however, Hutchins (1995, pp. 284–285) for possible exceptions to Vygotsky's view of the social origin of higher mental functions.

[24] See Bessy (2000) for a discussion of the different levels of analysis in Hutchins' work.

[25] Hutchins distances himself from an extreme situated or contextualised perspective in the introduction (p. 25), when he states that little would have been changed in his account had he made his observations on another Navy vessel.

[26] Much of the organisational behaviour literature concerned with this issue focuses on the hybrid forms of organisation that emerge in the context of the diffusion of organisational practices across national boundaries. For a survey, see Doeringer, Lorenz and Terkla (2002). For the 'societal effects' approach to national specificity in work organisation, see Maurice and Sorge (2000). For a survey of the industrial relations and business literatures dealing with national specificity in management practice within Europe, see Lane (1989).

[27] See Lam's (1996) useful comparison of the work practices of Japanese and British engineers in the case of an Anglo-Japanese joint venture in the electronics sector.

## References

Abegglen, J. and G. Stalk: 1985, *Kaisha: The Japanese Corporation* (New York: Basic Books).

Aoki, M.: 1988, *Information, Incentives and Bargaining in the Japanese Economy* (Cambridge, MA: Cambridge University Press).

Bessy, C.: 2001, "Is the Growth of the Firm Limited by Tacit Knowledge? The Evolutionary Approach to the Firm Revisited by the Distributed Cognition Perspective", Working Paper, Centre d'Etudes de l'Emploi, Noisy-le-Grand, France.

Blanning, R.: 1996, "Knowledge, Metaknowledge and Explanation in Intelligent Organisational Models", in M. Warglien and M. Masuch (eds.), *The Logic of Organisational Disorder* (Berlin: Walter and Gruyter).

Blumer, H.: 1969, *Symbolic Interactionism: Perpsective and Method* (Berkeley: University of California Press).

Brown, J.S., A. Collins and P. Duguid: 1989, "Situated Cognition and the Culture of Learning", *Education Researcher* 18: 32–42.

Brown, J.S. and P. Duguid: 1991, "Organisational Learning and Communities of Practice: Toward a Unified View of Working, Learning and Innovation", *Organisation Science* 2(1): 40–57.

Carley, K.: 1986, "Efficiency in a Garbage Can: Implications for Crisis Management", in J. March and R. Weissinger-Baylon (eds.), *Ambiguity and Command*, pp. 165–194.

Carley, K.: 1995, "Computational and Mathematical Organisational Theory: Perspectives and Directions", *Computational and Mathematical Organisational Theory* 1: 39–56.

Cohen, M.: 1998, "Foreword", in M. Preitula, K. Carley and Les Gasser (eds.), *Simulating Organisations: Computational Models of Institutions and Groups* (Cambridge, MA: MIT Press).

Cohen, M. and P. Bacdayan: 1994, "Organisational Routines are Stored as Procedural Memory: Evidence from a Laboratory Study", *Organisation Science* 5(4): 554–568.

Cohen, M., R. Burkhart, G. Dosi, M. Egidi, L. Marengo, M. Warglien and S. Winter: 1995, "Routines and Other Recurring Action Patterns of Organisations: Contemporary Research Issues", Santa Fe Working Paper, Santa Fe, NM.

Cohen, M., J. March and J. Olsen: 1972, "A Garbage Can Model of Organisational Choice", *Administrative Science Quarterly* 17(1): 1–25.

Cole, M. and Y. Engeström: 1993, "A Cultural-Historical Approach to Distributed Cognition", in G. Salomon (ed.), *Distributed Cognition: Psychological and Educational Considerations* (Cambridge: Cambridge University Press), pp. 1–46.

Clancey, W.J.: 1997, *Situated Cognition: On Human Knowledge and Computer Representation* (Cambridge: Cambridge University Press).

Cyert, R. and J. March: 1963, *A Behavioural Theory of the Firm* (New Jersey: Prentice-Hall).

Doeringer, P., E. Lorenz and D. Terkla: 2002, "National Differences in High Performance Management: Evidence From Japanese Multinationals", *Cambridge Journal of Economics*, forthcoming.

Dore, R.: 1973, *Japanese Factory, British Factory* (Berkeley: University of California Press).

Dosi, G.: 1991, "Information, Competences and the Firm", Working Paper, University of Rome 'La Sapienza', Rome.

Dosi, G., R. Nelson and S. Winter (eds.): 2000, *The Nature and Dynamics of Organisational Capabilities* (Oxford: Oxford University Press).

Egidi, M.: 1995, "Accountants and Termites", in M. Cohen, R. Brakhart, G. Dosi, M. Egidi, L. Marengo, M. Warglien and S. Winter (eds.), *Routines and Other Recurring Action Patterns of Organisations: Contemporary Research Issues*, Appendix A, Santa Fe Working Paper, Santa Fe, NM.

Egidi, M.: 1996, "Routines, Hierarchies of Problems, Procedural Behavior: Some Evidence from Experiments", in K. Arrow et al. (eds.), *The Rational Foundations of Economic Behavior* (London: MacMillan).

Fischer, B. and A. Strauss: 1978, "Interactionism", in T. Bottomore and R. Nisbet (eds.), *A History of Sociological Analysis*, Ch. 12 (New York: Basic Books), pp. 457–498.

Fox, S.: 2000, "Communities of Practice, Foucault and Actor-Network Theory", *Journal of Management Studies* 37(6): 853–867.

Garfinkel, H.: 1967, *Studies in Ethnomethodology* (Engelwood Cliffs, N.J.).

Hutchins, E.: [1986] 1997, "Mediation and Automatization", in M. Cole, Y. Engeström and O. Vasquez (eds.), *Mind, Culture and Activity: Seminal Papers from the Laboratory of Comparative Human Development* (Cambridge: Cambridge University Press), pp. 338–353.

Hutchins, E.: 1995, *Cognition in the Wild* (Cambridge, MA: MIT Press).

Laboratory of Comparative Cognitive Development: 1983, "Culture and Cognitive Development", in P.H. Mussen (ed.), *Handbook of Child Psychology*, Vol. 1 (W. Kessen, volume editor) (New York: John Wiley), pp. 318–356.

Lam, A.: 1996, "Engineers, Management and Work Organisation: A Comparative Analysis of Engineers' Work Roles in British and Japanese Electronics Firms", *Journal of Management Studies*.

Lane, C.: 1989, *Management and Labour in Europe* (Aldershot, Hants: Edward Elgar).

Lave, J.: [1981] 1997, "What's Special About Experiments as Contexts for Thinking", in M. Cole, Y. Engeström and O. Vasquez (eds.), *Mind, Culture and Activity: Seminal Papers from the Laboratory of Comparative Human Development* (Cambridge: Cambridge University Press), pp. 57–68.

Lave, J.: 1988, *Cognition in Practice: Mind Mathematics and Culture in Everyday Life* (Cambridge: Cambridge University Press).

Lave, J., M. Murtaugh and O. de la Rocha: 1984, "The Dialectic of Arithmetic in Grocery Shopping", in B. Rogoff and J. Lave (eds.), *Everyday Cognition: Development in Social Context* (Cambridge, MA: Harvard University Press), pp. 67–94.

Lave, J. and E. Wenger: 1991, *Situated Learning: Legitimate Peripheral Participation* (Cambridge: Cambridge University Press).

Lazaric, N. and P. Mangolte: 1999, "Routines in Theory and in Practice: A Criticism of the Cognitive Perspective", *Revista de Economia Contemporânea* 5: 7–35.

Leigh Star, S.: 1998, "Working Together: Symbolic Interactonism, Activity Theory, and Information Systems", in Y. Engeström and D. Middleton (eds.), *Cognition and Comunication at Work* (Cambridge: Cambridge university Press), pp. 296–318.

Leont'ev, A.: 1981, "The Problem of Activity in Psychology", in J. Wertsch (ed.), *The Concept of Activity in Soviet Psychology* (Armonk, NY: M.E. Sharpe).

March, J. and H. Simon: [1958] 1993, *Organisations* (Cambridge, MA: Blackwell).

Masuch, M.: 1990, *Organisation, Management and Expert Systems: Models of Automated Reasoning* (Berlin: Walter de Gruyter), pp. 38–67.

Masuch, M. and Lapotin: 1989, "Beyond Garbage Cans: An AI Model of Organisational Choice", *Administrative Science Quarterly* 34.

Maurice, M. and A. Sorge: 2000, *Embedding Organisations: Societal Analysis of Actors, Organisation and Socio-Economic Context* (Amsterdam: John Benjamins).

McDermott, R.: 1998, "Learning Across Teams: The Role of Communities of Practice", *Knowledge Management Review*, May/June.

Newell, A., P.S. Rosenbloom and J.E. Laird: 1989, "Symbolic Architectures for Cognition", in M.I. Posner (ed.), *Foundations of Cognitive Science* (Cambridge, MA: MIT Press).

Newell, A., J.C. Shaw and H. Simon: [1958] 1989, "Elements of a Theory of Human Problem Solving", *Models of Thought, Vol. II* (New Haven: Yale University Press), pp. 6–19.

Newell, A. and H. Simon: 1972, *Human Problem-Solving* (New Jersey: Prentice-Hall).

Newell, A. and H. Simon: 1976, "Computer Science as Empirical Enquiry: Symbol and Search", *Communications of the AMC*, 19(3).

Nightingale, P.: 2001, "If Nelson and Winter Are Only Half Right about Tacit Knowledge, Which Half? A Reply to David, Foray and Cowan", Paper presented at the DRUID Summer Conference, Aalborg, Denamark.

Nooteboom, B.: 2001, *Learning and Innovation in Organisations and Economies* (Oxford: Oxford University Press).

Padgett, J.: 1980, "Managing Garbage Can Hierarchies", *Administrative Science Quarterly* 25: 583–604.

Polayni, M.: 1969, *Knowing and Being: Essays by Michael Polayni* (London: Routledge and Kegan Paul).

Rogoff, B.: 1990, *Apprenticeship in Thinking: Cognitive Development in Social Context* (Oxford: Oxford University Press).

Rogoff, B. and J. Lave: 1984, *Everyday Cognition: Development in Social Context* (Cambridge, MA: Harvard University Press).

Simon, H.: 1998, *The Sciences of the Artificial*, 3rd edn. (Cambridge, MA: MIT Press).

Snyder, W.: 1998, "Communities of Practice: Combining Organisational Learning and Strategy Insights to Create a Bridge to the 21st Century", Document, Social Capital Group, Cambridge, MA.

Suchman, L.: 1987, *Plans and Situated Actions: The Problem of Human-Machine Communication* (Cambridge: Cambridge University Press).

Suchman, L.: 1998, "Constituting Shared Workspaces", in Y. Engeström and D. Middleton (eds.), *Cognition and Communication at Work* (Cambridge: Cambridge University Press), pp. 35–60.

Wenger, E.: 1998a, *Communities of Practice: Learning, Meaning and Identity* (Cambridge: Cambridge University Press).

Wenger, E.: 1998b, "Communities of Practice; Learning as a Social System", *Systems Thinker*, June.

Wertsch, J.V.: 1979, "Introduction", in J. Wertsch (ed.), *The Concept of Activity in Soviet Psychology* (Armonk, NY: M.E. Sharpe).

Wertsch, J.V.: [1987] 1997, "Collective Memory: Issues from a Sociohistorical Perspective", in M. Cole, Y. Engeström and O. Vasquez (eds.), *Mind, Culture and Activity: Seminal Papers from the Laboratory of Comparative Human Development* (Cambridge: Cambridge University Press), pp. 226–232.

Vygotsky, L.S.: 1979, "The Genesis of Higher Mental Functions", in J. Wertsch (ed.), *The Concept of Activity in Soviet Psychology* (Armonk, NY: M.E. Sharpe).