

Bridging Epistemologies: The Generative Dance Between Organizational Knowledge and Organizational Knowing

Scott D. N. Cook • John Seely Brown

Department of Philosophy, San Jose State University, San Jose, California 95192, and Xerox Palo Alto Research
Center, 3333 Coyote Hill, Palo Alto, California 94304, Cook@parc.xerox.com

Xerox Palo Alto Research Center, 3333 Coyote Hill, Palo Alto, California 94304, JSB@parc.xerox.com

It's funny what's happened to this word *knowing* . . . The actual *act* of apprehending, of making sense, of putting together, from what you have, the significance of where you are—this [now] oddly lacks any really reliable, commonly used verb in our language, . . . [one] meaning the *activity* of knowing . . . [Yet], every culture has not only its own set *body* of knowledge, but its own *ways* of [knowing].

Sir Geoffrey Vickers

Abstract

Much current work on organizational knowledge, intellectual capital, knowledge-creating organizations, knowledge work, and the like rests on a single, traditional understanding of the nature of knowledge. We call this understanding the “epistemology of possession,” since it treats knowledge as something people possess. Yet, this epistemology cannot account for the *knowing* found in individual and group *practice*. Knowing *as* action calls for an “epistemology of practice.” Moreover, the epistemology of possession tends to privilege explicit over tacit knowledge, and knowledge possessed by individuals over that possessed by groups. Current work on organizations is limited by this privileging and by the scant attention given to knowing in its own right. Organizations are better understood if explicit, tacit, individual and group knowledge are treated as four distinct and coequal forms of knowledge (each doing work the others cannot), and if knowledge and knowing are seen as mutually enabling (not competing). We hold that knowledge is a tool of knowing, that knowing is an aspect of our interaction with the social and physical world, and that the interplay of knowledge and knowing can generate new knowledge and new ways of knowing. We believe this *generative dance* between knowledge and knowing is a powerful source of organizational innovation. Harnessing this innovation calls for organizational and technological infrastructures that support the interplay of knowledge and knowing. Ultimately, these concepts make possible a more robust framing of such epistemologically-centered concerns as core competencies, the management of intellectual capital, etc. We explore these views through three brief case studies drawn from recent research.

(*Knowledge; Knowing; Epistemology; Practice*)

Introduction

In recent years, knowledge has become a prominent theme in the organizational literature. However, in such discussions, as in informal contexts, knowledge is typically spoken of as though it were all of a piece, as though essentially it comes in only one kind. It is our contention that there are, in fact, a number of distinct forms of knowledge, and that their differences are relevant, both theoretically and practically, to an effective understanding of organizations.

There is now much discussion of organizational knowledge, knowledge-based organizations, knowledge-creating organizations, knowledge work, etc. There are numerous related themes such as organizational learning, the collective mind (Weick and Roberts 1993), and the organizational brain. It has become common to talk of knowledge in the context of both individuals and groups, and even to consider knowledge in explicit and tacit senses (where, for example, explicit knowledge is treated as knowledge that can be spelled out or formalized, and tacit knowledge as that associated with skills or “know-how”). Accordingly, there are discussions about: how explicit knowledge acquired by individuals in an organization is associated with “learning” at the level of the

organization (March and Olsen 1976, Argyris and Schon 1978, Sims and Gioia 1986, Simon 1991, Sitkin 1992); how a group's mastering of explicit routines can be an aspect of organizational memory (Cohen and Bacdayan 1994); how the tacit skills of an individual can and cannot be tapped for the benefit of the organization (Nonaka 1994, Nonaka and Takeuchi 1995, Spender 1996); and how the activities of groups can constitute organizational learning (Weick and Westley 1996, Weick 1991). Meanwhile, such concepts are clearly vital to such concerns as the management of intellectual capital (Stewart 1997), core competencies (Hamel and Prahalad 1994), and innovation (Leonard-Barton 1995). Increasingly, such work has pushed provocatively and insightfully at the boundaries of the theoretical frames used in understanding knowledge and organizations—as in Weick and Roberts's (1993) application of “taking heed” and “mindfulness” to operations of teams; in Cohen and Bacdayan's (1994) use of notions of procedural memory from psychology as a way of understanding organizational routines; in what Hutchins (1991, p. 2) sees as the “pattern of communication” within the “cognitive system” of a group; in Nonaka and Takeuchi's spiral of knowledge creation; and in Kogut and Zander's (1996) considerations of the interplay between individuals' social knowledge and the organizing principles of work in explaining what organizations know how to do.

Yet, even in this growing body of literature that explores epistemologically significant themes, there typically remains an expressed or implied tendency to treat knowledge as being essentially of one kind. That is, the epistemology assumed in the literature tends to privilege the individual over the group, and the explicit over the tacit (as if, for example, explicit and tacit knowledge were two variations of one kind of knowledge, not separate, distinct forms of knowledge). The former tendency is reflected in the insistence that organizational learning is really about individual learning since “All learning takes place inside individual human heads . . .” (Simon, 1991 p. 125). The latter, meanwhile, can be seen in Nonaka's argument that “While tacit knowledge held by individuals may lie at the heart of the knowledge creating process, realizing the practical benefits of that knowledge centers on its externalization . . .” where “externalization” for Nonaka entails a process of “converting” tacit knowledge into explicit knowledge (1994, p. 20). Cohen and Bacdayan, meanwhile, contend that organizational routines arise when “individuals store components of a routine as a procedural memory” (1994, p. 554). And even Weick and Roberts have made the epistemologically provocative move of describing “collective mind” in terms

of “a distinct higher-order pattern of interrelated activities” grounded in and emerging from “individual actions” (1994, p. 374). Meanwhile Hutchins (1991, p. 284) speaks of investigating the “ways in which the cognitive properties of human groups may depend on the social organization of individual cognitive capabilities.”

As we will detail below, we believe that the tendency to treat all knowledge as being essentially the same severely limits the current work on epistemologically-relevant organizational themes, both theoretically and operationally. Theoretically, these tendencies fail to honor aspects of the distinction between explicit and tacit, and individual and group knowledge that we see as germane to understanding the acquisition, maintenance, and exercise of competencies by individuals and groups. Practically, it limits our ability to assess and support these competencies in their own right.

The first contention of this paper is that each of the four categories of knowledge inherent in the explicit/tacit and individual/group distinctions is a distinct form of knowledge on equal standing with the other three (i.e., none is subordinate to or made up out of any other). Also, this distinct character is reflected in the fact that each form of knowledge does work that the others cannot. We view these four forms of knowledge as constituting the appropriate focus of what we call *the epistemology of possession*, since these forms of “what is known” are typically treated as something people *possess*.¹ To say, for example, “Robert knows auto mechanics” points to Robert *possessing* knowledge of auto mechanics.

The second contention is that not all of what is known is captured by this understanding of knowledge. Put another way, there is more epistemic work being done in what we know how to do than can be accounted for solely in terms of the knowledge we possess.² So, in addition to talking about the four distinct forms of knowledge we *also* want to be able to speak about the epistemic work done by human action itself—that is, about what is *part of practice* as well as what is *possessed in the head*. To say, for example, “Robert is fixing cars” points not only to knowledge he possesses but also to things he is doing. To give an account of what Robert knows, we claim, calls for an understanding of the epistemic work done, which needs to include both the knowledge he possesses and the actions he carries out.

Borrowing from the epistemological perspective of the American Pragmatist philosophers, we call what is possessed “knowledge” and what is part of action “knowing.” Individuals and groups clearly make use of knowledge, both explicit and tacit, in what they do; but not everything they know how to do, we argue, is explicable solely in terms of the knowledge they possess. We believe

that understanding of the epistemological dimension of individual and group action requires us to speak about both knowledge *used in action* and knowing *as part of action*. Therefore, in addition to the traditional epistemology of possession, there needs to be, in our view, a parallel *epistemology of practice*, which takes ways of knowing as its focus. By this, we do not mean that practice needs to be brought under the umbrella of traditional epistemology (nor do we mean that all of human action needs to be accounted for epistemologically). Rather, we contend that there needs to be a radical expansion of what is considered epistemic in its own right, which includes knowledge and knowing.

Furthermore, we do not see knowledge and knowing as competing, but as complementary and mutually enabling (see Figure 1).³ Indeed, as we will spell out in detail in what follows, understanding what is entailed in bridging the two epistemologies provides a more robust account of such matters as: how individuals and groups can draw on tacit and explicit knowledge simultaneously; how what individuals know tacitly can be made useful to groups; and how explicit instructions can be made more useful aids for the development of tacit skills. Also (and quite importantly) we see the interplay of knowledge and knowing as a potentially generative phenomenon. That is, for human groups, the source of *new* knowledge and knowing lies in the use of knowledge as a tool of knowing within situated interaction with the social and physical world. It is this that we call the *generative dance*. Understanding the generative dance (how to recognize, support, and harness it) is essential, we believe, to understanding the types of learning, innovation, and effectiveness that are prime concerns for all epistemologically-oriented organizational theories.

In what follows, we explore the epistemologies of possession and practice and some implications of our perspective. We first sketch out our interpretation of the epistemology of possession, along with what we see as its


strengths and limitations. Then we offer what in our view are some essential elements of an epistemology of practice—in particular, we define what we mean by (1) the term *practice*, (2) the distinction between *knowledge* and *knowing*, (3) the Pragmatist philosopher John Dewey’s concept of *productive inquiry*, (4) the notion of *interaction with the world*, and (5) the idea of *dynamic affordance*. Following this, we look at how seeing knowledge as a tool of knowing can help explain how individuals and groups draw on all four forms of knowledge and, importantly, how the interplay between knowledge and knowing can generate new knowledge and new ways of knowing. In the final section, we explore these ideas in the context of three cases, and consider some broader implications of them for a more robust understanding of the epistemological dimension of organized human activity.

1. The Epistemology of Possession

Each of the four categories that come from the explicit/tacit and individual/group distinctions identifies a unique and irreducible form of knowledge. We see each of the four as on equal footing with the other three, and hold that no one of them can be derived from or changed into one of the others. We believe that each needs to be understood *conceptually* as distinct, in no small part because *in practice* each does work that the others cannot. In arguing for this position, we first address the conventional inclination to treat knowledge either as if it were all of a piece or, if different forms are considered, to privilege explicit over tacit and individual over group knowledge.

Privileging the explicit and the individual is not at all unique to organizational studies. It reflects the dominant epistemology of Western culture for the last three centuries, at least. This view is often referred to as the Cartesian view, given its substantial grounding in the work

Figure 1 Knowledge and Knowing



of the seventeenth-century French philosopher René Descartes. For Cartesians past and present, the individual, indeed the individual analytic thinker, is taken as primary. All knowledge, accordingly, is believed to be best acquired through reason and the use of concepts and methods that are freed as much as possible from the fallibilities of our senses or the exigencies of given situations.

Descartes' famous "Cogito ergo sum" (I think therefore I am) is both a beginning and a conclusion for the traditional epistemology. It is the conclusion that the thinking self is the one thing we cannot doubt—everything else, from the impressions of our senses to "objective" claims about the world, is subject to one or another degree of uncertainty. It is through analytic reasoning, Cartesians maintain, that we can best minimize or "control for" the clouding influences of our senses and subjective impressions, and thus acquire our most reliable knowledge about the world. It is a beginning in that the thinking (or reasoning or doubting) self becomes the one fundamental, irreducible starting point for any search for knowledge about the world, and the repository for that knowledge once acquired. All this should have a familiar ring to anyone who received a traditional introduction to "the scientific method" and "the scientific worldview."

What follows from all this has become part of the conventional understanding of knowledge in our culture: the idea that knowledge, particularly anything that might pass as rigorous knowledge, is something that is held in the head of an individual and is acquired, modeled, and expressed most accurately in the most objective and explicit terms possible. It is this Cartesian tradition, as well, that we see underlying such statements quoted above as "All learning takes place inside individual human heads . . ." (Simon 1991) and ". . . realizing the practical benefits of [tacit] knowledge centers on its externalization . . ." (Nonaka 1994).

Our aim here, it should be noted, is not to reject the Cartesian epistemology wholesale. Rather, we wish to critique some of its elements that we believe have made difficult the development of a productive understanding of the forms of knowledge suggested by categories other than individual/explicit. We believe Cartesian epistemology needs to be broadened into an "epistemology of possession" that can incorporate a conceptually sound and useful understanding of knowledge possessed tacitly and knowledge possessed by groups.

Explicit/Tacit

The grip that the Cartesian tradition has had on the exploration of explicit and tacit knowledge has been particularly strong. When the idea of tacit knowledge is addressed, for example, it is most often treated as an

informal, inchoate, or obscure kind of knowledge, whose very nature calls for it to be made explicit in order to be truly understood or useful in practice. Indeed, the very term "tacit" suggests to many people (quite understandably) the sense that any such knowledge must be "hidden" from our understanding or "inaccessible" for practical purposes. We believe that this predilection of the traditional epistemology has held back the development of an understanding of the explicit/tacit distinction that is called for and increasingly needed, given the growth of significant work on epistemological themes in the literatures concerned with organized human action. Indeed, we base our claim that the explicit/tacit distinction is one between two separate forms of knowledge on practical utility: we argue that the distinction needs to be *conceptually* clear because, in practice, each form of knowledge does work the other cannot. A sounder, more robust conceptual understanding of the distinction should help make it possible to recognize, support, and harness the different forms of work that each, in fact, makes possible in practice.⁴

We base our understanding of the tacit/explicit distinction on the work of the scientist and philosopher Michael Polanyi (1983). Polanyi's distinction is exemplified very compellingly in the simple but rich example of riding a bicycle. Many people who say they can ride a bicycle will claim, when asked, that they do not know which way to turn the handlebars to prevent a fall to the left or right. However, since staying upright is part of knowing how to ride a bicycle, anyone who can ride must, by definition, know which way to turn the handlebars to avoid a fall. What they can't do is *say* which way to turn. So there's something known by everyone who can ride that most cannot say. What they can say in an example of what Polanyi called the explicit dimension of knowledge, while what is known by everyone who can keep upright on a bike is what he called the tacit dimension of knowledge.

Building on Polanyi, we argue that explicit and tacit are two distinct forms of knowledge (i.e., neither is a variant of the other); that each does work the other cannot; and that one form cannot be made out of or changed into the other. We explore these and other aspects of the distinction below, again beginning with the example of bicycle riding.

To be able to ride a bicycle, one needs to have the (tacit) knowledge of how to stay upright. This is knowledge one possesses; it is *not* the activity of riding itself but knowledge used in riding (you still possess the tacit knowledge even when you are not riding). Possessing this tacit knowledge makes it possible to keep upright, which is something that the explicit knowledge of which way to turn cannot do. We can't put a novice on a bicycle saying

“OK, take off—and if you start to fall like so, turn this way” and expect the person to be able to ride successfully. The novice would have the explicit knowledge but not the necessary tacit knowledge. Whatever epistemic work that explicit bit of knowledge can make possible, it cannot do *all* of the work that is necessary for someone to know how to ride. In order to acquire the tacit knowledge, a novice has to spend a certain amount of time on a bicycle. Indeed, it would even be possible for someone to be able to say in great technical detail what must be done to keep a bicycle upright, yet still be unable to ride one. No amount of explicit knowledge alone can enable someone to ride; it simply cannot enable all the necessary epistemic work.

At the same time, we argue that each form of knowledge can often be used as an *aid* in acquiring the other. If you know how to ride, for example, you might use your tacit knowledge to ride around in a way that helps you discover which way you turn when you begin to fall. Likewise, if a novice is told how to turn to avoid a fall, that explicit knowledge could be used while learning to ride as an aid in getting a feel for staying upright. However, neither tacit nor explicit knowledge can be used by itself to acquire the other: one must also, at the very least, get on a bicycle (an important point, to which we will return shortly).

We can now see that each form of knowledge does its own work. Explicit knowledge can be used as an aid to help acquire the tacit knowledge, but cannot by itself enable one to ride. The tacit knowledge is necessary in being able to ride, but it does not by itself enable a rider to say which way to turn.

Furthermore, it is important not to mistake using one form of knowledge as an aid in acquiring the other with one form being “converted” into the other. Tacit knowledge cannot be turned into explicit, nor can explicit knowledge be turned into tacit. If you ride around using your tacit knowledge as an aid to discovering which way you turn, when you ultimately acquire the explicit knowledge you still possess the tacit knowledge, and you still use it in keeping upright. When we ride around with the aim of acquiring the explicit knowledge, we are not performing an operation on our tacit knowledge that turns it into explicit knowledge; we are using the tacit, within the activity of riding, to generate the explicit knowledge. The explicit knowledge was not lying inside the tacit knowledge in a dormant, inchoate, or hidden form; it was generated in the context of riding with the aid of what we knew tacitly. Likewise, if you know explicitly which way to turn but cannot ride, there is no operation you can perform on that explicit knowledge that will turn it into

the tacit knowledge necessary to riding. That tacit knowledge is acquired on its own; it is not made out of explicit knowledge. Prior to being generated, one form of knowledge does not lie hidden in the other.

Also, there is no guarantee that one form will always be a useful aid to acquiring the other. In fact, in some cases using one can be a hindrance to acquiring the other. In learning how to drive, for example, you may be told (explicitly) to accelerate when coming out of a turn, only to be told later that you are using this knowledge mechanically “as a crutch” rather than “getting a feel for it.” Similarly, in learning a skill like dancing or tennis many people experience a period when explicit knowledge about how to move one’s feet or hold one’s shoulders can actually impair one’s ability to acquire the tacit knowledge necessary to performing the skill in a fluid or masterful way. Even experts in a given skill can find their ability to use their tacit knowledge “thrown off” when they are asked to describe explicitly what they are doing.

Individual/Group

We have also inherited a cultural predilection for privileging the individual over the group. Whether stated emphatically or present implicitly, a sense that whatever can be said about groups actually “boils down” to things about individuals is taken almost as though it were self-evident, and particularly so when the concern at hand is an epistemological one (Cook 1994). As the Cartesian view would have it, it is the *individual* thinker who is the primary (if not exclusive) wielder and repository of what is known. This predilection is reflected, for example, in Simon’s insistence (noted above) that all learning takes place inside the heads of individuals. For many who are not as orthodox as Simon, such topics as “organizational learning,” “organizational knowledge,” or “organizational routines” are still spoken of in ways that often leave it unclear as to whether groups are being treated on an equal footing with individuals or as a derivative of them. (This is often so, it should be noted, even in cases where it is not authors’ intention either to address or to dodge the issue.)

In recent years, however, there has been a growing volume of research and publication that has begun to treat groups and organizations in their own right. This has been an implicit concern in our own work as well as that of a number of our colleagues at Xerox PARC and the Institute for Research on Learning. This trend is also strongly suggested in the literature treating such concepts as “communities of practice” (Wenger 1997, Brown and Duguid 1991), “core competencies” (Hamel and Prahalad 1994), “situated cognition,” “legitimate peripheral participation” (Lave and Wenger 1991), and the “spiral of organizational knowledge creation” (Nonaka and Takeuchi 1995).

Discussions of communities of practice look at how individuals establish themselves and function as a group by engaging in practices that are unique to or characteristic of that group. Within the growing body of work on core competencies one can see serious attention being given to how teams, as well as individuals, do “real work” and how that work can be supported, enriched, and directed. The concept of legitimate peripheral participation, originally used to explore apprenticeship learning, takes as its central concern the role of participation by seemingly peripheral individuals in the innovative and very *central* capacities of the group itself. In more and more instances, authors are addressing such epistemological issues at the level of the group, including recent direct explorations of such terms as “organizational knowledge” and “organizational epistemology” (von Krogh and Roos 1995). By taking the group as a primary unit of analysis, such approaches, implicitly at least, treat groups as something to be investigated in their own right with respect to epistemological concerns.

As with the explicit/tacit distinction, we propose that individuals and groups each do epistemic work that the other cannot. So, for example, while only individual physicians know how to diagnose nephritis using palpation (groups do not have hands), the knowledge of what constitutes acceptable and unacceptable practice in nephrology is possessed by nephrologists as a group. Likewise, while individual copier technicians have a sense of how a particular copier ought to sound when operating properly (groups do not have ears), it is a group of technicians that possess “war stories” about what odd noises can mean. Indeed, an individual technician’s account only becomes a “war story” when it is held in common and can be used by the group in its discussions about machines (Orr 1996). In both cases, part of what is known about a given domain is possessed by individuals, part by groups. Individual technicians and nephrologists possess various bits of knowledge in their respective fields, but the “body of knowledge” of copier repair or nephrology is possessed by groups, not by individuals. Put another way, the body of knowledge of a group is “held in common” by the group. We do not expect every individual in a group (discipline, profession, craft, etc.) to possess everything that is in the “body of knowledge” of that group (in fact, this is likely to be impossible, unnecessary, and perhaps even undesirable). The body of knowledge is possessed by the group as a whole and is drawn on in its actions, just as knowledge possessed by an individual is drawn on in his or her actions. The work done by a group, as informed by the body of knowledge it possesses, is work that is epistemically distinct from work done by an individual in it, as informed by the knowledge he or she possesses.⁵

With respect to both distinctions, the lesson we wish to draw here is *not* that we ought now to reverse tradition and privilege the group and the tacit over the individual and the explicit. Indeed, our aim has been to argue for an expanded epistemology of possession that includes each of four types of knowledge and treats each as distinct from (not superior to) the other three, both conceptually and in the sense of each doing work that the others cannot.

2. Toward an Epistemology of Practice

We are now able to focus on an important aspect of what people know that is *not* captured by the four forms of knowledge considered above. In the bicycle example we argued that tacit and explicit knowledge alone are insufficient in acquiring the ability to ride; what has to be added is the actual act of riding (or trying to). This leads us now to make a specific claim: *the act of riding a bicycle does distinct epistemic work of its own*. Indeed, we hold that this type of epistemic work is an inextricable facet of human action itself, not something people possess. We mark this distinction by referring to it as “knowing” rather than “knowledge.” Furthermore, we believe that knowing does not belong to an epistemology of possession, but rather that it calls for an epistemology of practice. Following Vickers’ (1976, p. 2) assertion that every human group “has not only its own set *body* of knowledge, but its own *ways* of [knowing],” we now turn to outlining some of what we believe “knowing” and an “epistemology of practice” entail. In particular, we propose specific understandings of (1) the term *practice*, (2) the distinction, drawn from the Pragmatists, between *knowledge* and *knowing*, (3) John Dewey’s concept of *productive inquiry*, (4) the notion of *interaction with the world*, and (5) the idea of *dynamic affordance*.

Practice

Practice implies doing. Intuitively, it refers to things we do as individuals and as groups. Conceptually, practice has received a growing amount of careful theoretical attention in recent years (see, for example, Bourdieu 1977, Turner 1994). In common usage, “practice” can mean either to develop a competency through drill or rote actions as in “to practice the piano” or to exercise a competency as in “to practice medicine.” The former suggests drill in preparation for doing the “real work,” while the latter suggests the “real work” itself. In our use of the term, we mean doing real work: the practice of engineers, managers, physicians, woodworkers, etc. (in which, meanwhile, drill and other rotelike activities can play an important part).

For our purposes, then, we intend the term “practice” to refer to *the coordinated activities of individuals and*

groups in doing their “real work” as it is informed by a particular organizational or group context. In this sense, we wish to distinguish practice from both behavior and action. Doing of any sort we call “behavior,” while “action” we see as behavior imbued with meaning. By “practice,” then, we refer to action informed by meaning drawn from a particular group context. In the simplest case, if Vance’s knee jerks, that is behavior. When Vance raps his knee with a physician’s hammer to check his reflexes, it is behavior that has meaning, and thus is what we call action. If his physician raps his knee as part of an exam, it is practice. This is because the meaning of her action comes from the organized contexts of her training and ongoing work in medicine (where it can draw on, contribute to, and be evaluated in the work of others in her field).

Knowledge and Knowing

Drawing a distinction between knowledge and knowing may seem at first pass an unduly subtle point. We believe it is at root quite a substantial one, both epistemologically and in its implications for understanding organized human activity. Above, we have expanded our understanding of knowledge to include the forms suggested by the explicit/tacit and individual/group distinctions. With respect to all four forms, however, we have maintained the sense of knowledge as something that is possessed. When we say “Miriam has knowledge of physics,” the knowledge is something that Miriam possesses (as concepts, rules, procedures, etc.). Furthermore, her knowledge (whether explicit or tacit) is abstract since it is something that is *about* but not *in* the tangible world. And it is static, in that possessing it does not require that it be always in use: When Miriam is playing tennis or sleeping she still has knowledge of physics. Finally, while knowledge itself is static, it is common to see it as necessary to action: “Miriam can solve the problem because she has knowledge of physics” or “Miriam cannot solve the problem until she acquires knowledge of the conservation of angular momentum.” That is, knowledge is commonly thought of as something we *use* in action but it is not understood to *be* action.

Accordingly, we use the term “knowing” to refer to the epistemological dimension of action itself. By “knowing” we do not mean something that is *used in* action or something *necessary to* action, but rather something that is a *part of* action (both individual and group action). “Knowing” refers to the epistemic work that is done as part of action or practice, like that done in the actual riding of a bicycle or the actual making of a medical diagnosis. Knowing is dynamic, concrete, and relational. If we talk about André reflecting “knowing” in physics, our focus

is on what he is actually doing; it is on the ways he deploys the knowledge he possesses in his interactions with the materials of a specific concrete task in physics (such as testing an experimental laser design).

In developing an understanding of the knowledge/knowing distinction, we have found it useful to draw on the work of the American philosophical school of Pragmatism, in particular the work of John Dewey, as an alternative to the dominant Cartesian perspective. Those interested in organizations have generally seen the work of the Pragmatists as limited essentially to educational settings. We believe that a new look at the Pragmatist perspective can yield very important and timely implications for organizations of all sorts. The recent resurgence of interest in American Pragmatism, which has centered on Dewey (see, for example: Hickman 1990, Rorty 1982), makes the reexamination of this perspective even more timely for organizational concerns.

A basic conviction of the Pragmatist perspective in both theory and practice is that our primary focus should *not* be (solely) on the likes of abstract concepts and principles (as has been common more broadly in philosophy and the social sciences) but on concrete action. Pragmatists have been centrally concerned with doing, particularly forms of doing that entail making or producing something (from technologies to ideas). Accordingly, when it comes to questions of what we know and how we know, the Pragmatist perspective takes a primary concern not with “knowledge,” which is seen as abstract and static, but with “knowing,” which is understood as part of concrete, dynamic human action. Following the Pragmatist perspective, for us “knowing something” refers to an *aspect of* action, not to something assumed to underlie, enable, or be used in action.⁶ By “knowing” we mean that aspect of action or practice that does epistemic work.

“Knowing,” Dewey maintained, “is literally something which we do,” not something that we possess. For Dewey, to talk about activity in terms of knowledge is to mistake an abstract, static concept for a concrete, dynamic activity. It is to make a kind of category error. To be accomplished in a profession, discipline, or craft, for example, is necessarily tied up with practicing it. This does not mean that its body of knowledge is useless to practice, only that it is not the same as the epistemic dimension of practice. An accomplished engineer may possess a great deal of sophisticated knowledge; but there are plenty of people who possess such knowledge yet do not excel as engineers (as is often observed in many fields). This means that if you want to understand the essentials of what accomplished engineers know, you need to look at what they do as well as at what they possess. It also means that our fundamental understanding of the relationship

between a body of knowledge and activities of a practice must change: we must see knowledge as *a tool at the service of knowing* not as something that, once possessed, is all that is needed to enable action or practice. (Improved practice may not always be the product of acquiring more knowledge; at times it may be the result of developing innovative ways of using knowledge already possessed.)

This Pragmatist focus on action has broad implications for those areas where organizational and epistemological concerns intersect. And the value of these implications can be carried further, we believe, by drawing on the key Deweyan concept of “productive inquiry.”

Productive Inquiry

One of the most important things that knowing can do in using knowledge as a tool is what Dewey called “productive inquiry.” To engage in productive inquiry is to be actively pursuing a problem, puzzle, point of fascination, object of wonder, or the like; it is to seek an answer, solution, or resolution. It is *inquiry* because what motivates us to action is in some sense a query: a problem, a question, a provocative insight, or a troublesome situation. It is *productive* because it aims to produce (to make) an answer, solution, or resolution. Productive inquiry includes a broad range of actions from the problem solving of mathematics to computer programming to fixing a photocopier to finding the proper placement of the voice in singing. *Productive inquiry is that aspect of any activity where we are deliberately (though not always consciously) seeking what we need, in order to do what we want to do.*

Productive inquiry is not a haphazard, random search; it is informed or “disciplined” by the use of theories, rules of thumb, concepts, and the like. These tools of productive inquiry are prime examples of what Dewey understands the term “knowledge” to mean. Conversely, using knowledge in this way is an example of that particular form of knowing that Dewey called “productive inquiry.” So, using “knowledge” in productive inquiry gives inquiry a systematic or disciplined character: just as knowledge is a tool of knowing, so must knowing respect the demands and constraints of knowledge. (To wield any tool skillfully, we must respect the constraints it places on our actions in using it, as the haphazard use of a hammer can all too painfully demonstrate.)

Significantly, Dewey also saw knowledge as one of the possible outcomes of productive inquiry: one end result of engaging in the (situated, dynamic) activity of productive inquiry is the production of (abstract, static) knowledge, which then can be used as a tool of further knowing, including knowing in the mode of productive inquiry.

Building on these key points from Dewey, we make a number of further arguments about the distinction between knowledge and knowing. Knowledge by itself cannot enable knowing. As a tool, knowledge disciplines knowing, but does not enable it any more than possession of a hammer enables its skillful use. Likewise, the principles of engineering alone cannot enable an accomplished engineer to engage in the productive inquiry of resolving a difficult design problem. However, it is precisely such things as the principles of engineering that an accomplished engineer uses *in practice* as tools in addressing a problem at hand, in interacting with it through the use of those tools, in seeking to resolve a design problem.

Furthermore knowing should not be confused with “tacit knowledge.” As we have defined tacit knowledge, it is a tool or an aid to action, not part of action itself. Everyone who can ride a bike can be said to know tacitly which way to turn to avoid a fall, whether or not they are at that moment actually riding. Knowing requires present activity. Tacit knowledge does not. Knowing makes use of tacit knowledge as a tool for action—as when we ride around on a bike using our tacit knowledge to stay upright (acquiring the tacit knowledge of how to stay upright, meanwhile, is acquiring know-how useful to bike riding.) Finally, tacit knowledge alone does not enable us to ride; there is more epistemic work that needs to be done. Being able to ride requires interaction between the (tacit) knowledge we possess and the present activity of being in motion on a bike. The activity of riding, itself, is a form of knowing; it does distinct epistemic work. *Knowing is that aspect of action (or practice) that does epistemic work—including doing things we know how to do, and (through productive inquiry) producing what we need, in order to do something we want to do, which can include producing new knowledge.* We will explore this notion further in the next two sections.

Interaction with the World

We act within the social and physical world, and since knowing is an aspect of action, it is about interaction with that world. When we act, we either give shape to the physical world or we affect the social world or both. Thus, “knowing” does not focus on what we possess in our heads; it focuses on our interactions with the things of the social and physical world.

“Knowledge” is about possession; it is a term of predication. In all its forms we use it to indicate something an individual or group possesses, can possess, or needs to possess. “Knowing” is about relation: it is about interaction between the knower(s) and the world.

To interact with the world effectively we need to honor

it. One cannot make reliable objects through the haphazard use of clay or steel: it is possible to make the walls of a pot too thin or the span of a bridge too long: objects give way when design pushes them beyond the constraints of their materials. To make use of the power of materials, their inherent constraints must be honored. The master of a craft—whether potter or materials engineer—is constantly engrossed in a kind of conversation with the materials of his or her craft. The master puts out ideas by giving shape to the material, and “hears back” from it as he or she discovers and explores what the material can and cannot make possible. Part of what it means to master any craft is to learn how to turn the constraints of its materials into opportunities for design.

Similarly, in the social world, one must honor the strengths, limitations, and character of individuals and groups to engender coordinated and directed action or practice—as all good managers, football coaches, and orchestra conductors know, at least intuitively (as do the members of such groups).

Knowledge also helps us “honor” the world in our interactions with it. As noted above, knowing as an aspect of action can make use of bits of knowledge (in any of its forms) as tools. In doing so, the knowledge about the social and physical world “disciplines” our interaction with the world, just as the use of a pair of pliers gives particular form to how we interact with a bolt.

Within the relational and interactive character of knowing, the world shapes our actions by requiring that we honor it, just as we shape the world by interacting with it in a disciplined way. *Knowing is to interact with and honor the world using knowledge as a tool.* We will look more precisely at how this works in the next section.

Dynamic Affordance

We now wish to focus on some specific characteristics of “interaction with the world” that are at the center of our understanding of “knowing.” In doing so, we first borrow two general points from the work of the Spanish philosopher José Ortega y Gasset that frame “interaction with the world” in a way that further develops an alternative to the Cartesian frame of the “thinking self.” Then we explore the idea of “affordance,” as introduced in the work on perception by J. J. Gibson (1979) and as significantly developed in the design work of W. W. Gaver (1991, 1996). Finally we argue for our sense of what we call “dynamic affordance.”

Interaction with the social and physical worlds is a central concern in the work of Ortega. Very much in keeping with the American Pragmatists, Ortega abandoned the frame of the abstracted, analytic thinking self and throughout his work approached questions of epistemology, action, etc. from the perspective of “myself within

this context.” For Ortega, what we can know and what we can do are not discovered through an abstract Cartesian thought experiment, but are products of ongoing concrete interaction between “myself” (or “ourselves”) and the specifics of the social and physical “context” or “circumstances” we are in at any given time. “I invent projects of being and of doing,” Ortega insisted, “in light of circumstance” (1961a, p. 202).


In keeping with this, Ortega argues that in interacting with the world we encounter both “facilities” and “frustrations” (1961b). It is important to note that facilities and frustrations are *not* properties of the world, but properties that lie solely in our interaction with the world. The tensile strength of clay is a property of the world, but it becomes a facility or a frustration only when we are interacting with it (e.g., when we are making pots). Likewise, the bits of knowledge that members of a team may possess are a property of that social world. They can only become facilities or frustrations, however, when we are interacting with the group within the context of a specific piece of work (or when the members of the group interact with each other in such a context).

The phenomenon of certain properties arising solely in the context of interaction with the world can also be seen in connection with the idea of “affordance.” There is a common meaning of “affordance” that is a progenitor of the sense we have in mind, but it is one we need to go beyond, because it suggests a static (i.e., not “interactive”) character. This is the elemental sense of how a material, design, or situation “affords” doing something: metal affords making buckets; buckets afford carrying water; bucket brigades afford fire fighting.

This sense of affordance is reflected in everyday objects in ways that can attract a great deal of conscious attention or none at all. This is particularly true of objects that are the product of human design. What they afford can give rise to shape and fluidity or incoherence and clumsiness in our activities. This can be seen, for example, even in the simple case of an ordinary book. The design of a book, as distinct from a newspaper or a scroll, affords such things as skimming or random access by using a thumb index or flipping from one part of the text to another and back again.

A doorknob, to take another example, affords opening and closing a door. The particular design of a doorknob can afford fluid or clumsy action. In Figure 2 we show the design of a doorknob that affords pushing or pulling the door from the appropriate side. On the side where the door needs to be pushed, the knob is a flattened hemisphere flush with the door; it is a knob that would, in fact, be difficult to pull. On the opposite side the same shape is raised from the surface of the door and one’s fingers

Figure 2 Affordance



can fold easily around the edge so one's hand is almost invited to pull (particularly when paired with resistance from the door, if one should try pushing from that side). Although the design elements of common objects like books and door knobs are often at the border of our attention, they nonetheless can constitute important resources in our interactions with them (Brown and Duguid 1994).

How characteristics of the world give clues to our perceptions as to what we can and can't do with them is the sense of "affordance" that is explored in depth in the work of Gibson (1979). Gaver has carried this notion further by arguing for an understanding of affordance that is not primarily about perception but about relationships between characteristics of the world and issues of inherent concern to people. For Gaver (1991, 1996), questions of affordance with respect to elevation in architecture, for example, emerge as issues of "accessibility," which come from the relationship between elevation and the necessity of expending energy climbing to higher surfaces of support.

As we have indicated, there is a sense of affordance that lies beyond these inherently static senses, which deserves to be understood in its own right. We call this additional sense "dynamic affordance" and mean by it forms of affordance that emerge as part of the (dynamic) interaction with the world. In talking about design elements of ordinary objects, for example, we said that they "can give rise to shape and fluidity or incoherence and clumsiness in our activities." We would note now that "shape, fluidity, incoherence, and clumsiness" are not properties of the objects (i.e., of the world). Rather, like Ortega's facilities and frustrations, they are properties of our interactions with those objects.⁷ The emergence of these properties raises the question as to how we might deal with them: what use might we make of shape and fluidity, and how might we address incoherence and clumsiness are questions about what those properties of

interaction afford. They are questions about dynamic affordance.

What we mean by "dynamic affordance" has both an *intuitive* sense and a very particular *conceptual* sense. Both senses can be seen in the bicycle riding example. Intuitively, most of us understand that learning to ride requires "getting a feel" for what it is like to stay in balance, and we recognize that we need to get on a bike to acquire that knowledge. So, the activity of riding around *dynamically affords* the acquisition of the needed knowledge.

Conceptually, we see "dynamic affordance" as lying in the real and subtle interaction between the rider and the bike in motion. When bicycle wheels turn, they become gyroscopes—and like all gyroscopes their tendency is to remain in the plane of rotation: to get spinning bicycle wheels to tip to one side or the other requires that a force be applied to them that will overcome this gyroscopic tendency. A rider uses his or her body weight as that force: shifting one's weight pushes against the gyroscopic force of the moving wheels. This is what we do (or part of it) when we are riding or learning to ride. In the activity of riding, shifting our weight against the gyroscopic force of the wheels "dynamically affords" learning to stay upright; it also "dynamically affords" the enactment of that skill once acquired. These are things we can learn and do *only* when we are in dynamic interaction with bicycle wheels in motion. Without the dynamic affordance of that interaction there is no learning and no enactment of what is learned. Both are always inextricably tied to riding itself: without the activity of riding there is no gyroscopic force to be used or pushed against. This dynamic character is an essential element of our conceptual sense of "dynamic affordance."

Finally, because interaction between rider and bicycle dynamically affords *both* the acquisition of knowledge *and* the use of knowledge once acquired, we see it as doing epistemic work that the knowledge alone cannot. Indeed, we argue that dynamic affordance is intimately connected to the distinct epistemological form we have called "knowing." Dynamic interaction with the world opens the unique realm in which knowing takes place: the activity of addressing facilities and frustrations dynamically affords knowing.

We hold that dynamic affordance and knowing play an essential role in how knowledge—explicit and tacit, individual and group—is generated, transferred, and used in organizations. We also hold that these activities acquire particular shape and meaning from their organizational contexts—that is, they are not only actions; they are also practices. Consequently, understanding how what is known functions in organizations requires understanding


the interplay between the epistemology of possession and the epistemology of practice. It is to these matters that we now turn our attention.

3. Bridging Epistemologies

The four distinct forms of knowledge of the epistemology of possession as discussed above are displayed in the following figure.

The cells of Figure 3 array knowledge among the categories of individual/group and explicit/tacit. The upper left cell contains things an individual can know, learn, and express explicitly. Examples of things that would fit this cell would include (but certainly not be limited to) concepts, rules, and equations that typically are presented explicitly and are typically known and used by individuals. In the upper right are things that are also expressed explicitly yet typically are used, expressed, or transferred in a group. This includes, for example, stories about how work is done or about famous successes or failures (Orr 1990, 1996), as well as the use of metaphors or phrases that have useful meaning within a specific group. In the lower left are examples of tacit knowledge possessed by individuals, such as a skill in making use of concepts, rules, and equations or a “feel” for the proper use of a tool or for keeping upright on a bike. Finally, in the lower right is tacit knowledge possessed by groups. Although everyone has daily experience with this form of knowledge, it is perhaps the most difficult of the four to define. A working definition of it, however, is crucial to understanding the relationships among the four forms of knowledge and to appreciating the distinction between knowledge and knowing. We wish to label this form of knowledge with an expanded definition of the term “genre.”

Figure 3 Four Forms of Knowledge



Conventionally, “genre” is most familiar as a literary term, where it refers to types of literature—e.g., “novel” and “biography” are two distinct literary genres. Such genres do more than constitute a tidy scheme of classification: they also provide frames for understanding and interpreting what we read, without which a text could be utterly baffling or dangerously misleading. We read or “take in” a text one way if we understand it to be a novel, quite another if we think it is a biography. Importantly, it is the meaning of the term “novel” or “biography” that constitutes the genre, not the actual text or the meaning the text acquires when it is understood to belong to a given genre. As literary historians would remind us, this meaning is constantly evolving and undergoing a kind of implicit negotiation among writers, readers, and publishers as they read and discuss texts.

The power of genres to enable us to make sense of and use a text is so common in experience that we often are unconscious of it (Brown and Duguid 1994). The characteristics of the genre “newspaper” (folds, pulp paper, narrow columns of text, headlines, bylines, etc.) have meanings that we pay little, if any, conscious attention to; however, our ability to make sense of what newspapers say is highly dependent upon them. Without having been taught it or even reflecting on it consciously, most of us “read” the importance of front-page stories that appear above a newspaper’s fold as greater than those that appear below it.

Genres are no less important to the organizational world than they are to the literary world (Orlikowski and Yates 1994). A message from a coworker can signal one thing if it arrives as a handwritten note, but quite another if it is a printed memo or a formal letter. The genre (note, memo, or letter) provides a frame for interpreting a given text. Each of these forms of communication has a meaning understood and used by members of the organization. Indeed, employing genres is one way people in organizations communicate. As such, organizational genres acquire their very distinct (and quite effective) meanings not by deliberate design but (like that of “novel” and “biography”) in the course of their being used (or misused) in the context of work practices.

The power of organizational genres is reflected, for example, in the case of the manager who reads e-mail only as printed-out hard copy. After reading one such message, he phoned its author to tell him in no uncertain terms that such subjects “should never be circulated in a memo.” The author replied that he had “never written a memo like that,” and that he had discussed the subject with people “only through e-mail.” In their organization, memos and e-mail had in practice become two distinct genres; they

had acquired two distinct meanings (with which the manager was perhaps not yet familiar). What was appropriate to communicate in one genre was inappropriate in the other. The boss misread the author's message (not necessarily his words) because he took what was intended as one genre (one form of communication) to be another.

We wish to generalize this sense of "organizational genre" in defining what we mean by tacit group knowledge. For our purposes, "organizational genre" applies not only the distinctive and useful meanings a given group attaches to its various literary artifacts. It also applies to its various physical and social artifacts—that is, to different types of things (technologies or products, for example) and to different types of activities (such as ways of doing a task or types of meetings). These genres are not explicitly learned or known (although they can, for example, have explicit counterparts such as a label or a name). Their meanings emerge and undergo constant confirmation and/or modification through a kind of "negotiation in practice" as they are used in the context of the group's ongoing "real work." What an organizational genre means at any one time is, in a sense, the accretion or product of the history of its use: it is meaning laid down in past use, and tapped into or "reevoked" each time the members of the group use it in subsequent work. Accordingly, organizational genres have useful meaning solely in the context of a given group's practices—in this sense, they are possessed or "held in common" by that group and are unique to it.

Two organizations, for example, could have ad hoc workgroup meetings, in each case called "gatherings," that to an outsider could appear to be a single kind of semiformal update. However, the meaning that "gathering" has within each organization could be immensely different from its meaning in the other. In one, a "gathering" could be understood by that organization's members to be where "the real decisions" are made. In the other, it could be seen as a time to make subtle political moves. The events are alike. The names are the same. The genres are different. In each case, what "gathering" means is known by the members of that organization; it is group knowledge. And that knowledge can be used effectively or ineffectively (as were "e-mail" and "memo" in the above example) without any explicit discussion ever occurring. Accordingly, it is also tacit knowledge. For our purposes, then, this expanded sense of genre defines what we mean by group/tacit knowledge.

As group/tacit knowledge, genres do epistemically distinct work. This is reflected in a corporate executive's remarks on how a group of senior managers has made use of their organization's mission statement. "The senior staff developed the statement," he reported, "and the

group has a sense of what it means, and we make use of that meaning in our discussions." The group's "sense" of what the mission statement means does not refer to its text but to the mission statement itself. Like "novel" or "memo" or "gathering," it has become a genre within that group; it has acquired, in practice, tacit meaning that is known by the group. It can be used appropriately or inappropriately, effectively or ineffectively, but only in the context of group practice: as tacit/group knowledge (as an organizational genre), "mission statement" does the epistemically distinct work of giving shape and direction to the group's discussions. This is underscored by the executive's next remark. "But when I think about the statement on my own," he reflected, "it can . . . lead my thinking in directions I wouldn't go if I were working on the same issues along with members of the group." How the genre functions within group practice is distinct from its role in the executive's thinking on his own. The group's "sense" of what the mission statement means exemplifies what we have in mind by tacit knowledge possessed (or held in common) by a group.

Adding Knowing to Knowledge

Individuals and groups make use of knowledge in interaction with the things and activities of the social and physical world. Knowledge, as we have said, gives particular shape, meaning, and discipline to our interactions with the world. At the same time, it has been our contention that not all of what we know in interacting with the world lies in our knowledge: some also lies in our actions themselves. Riding a bicycle requires that we use tacit knowledge in interaction with a bicycle in motion: some of what we know in being able to ride is in that interaction itself. For the manager mentioned above, being able to have effective communication with his colleagues required using the right genre ("e-mail" rather than "memo") in his interactions with messages (the action of interpreting them) and their authors (the action of conversing with them): some of what he knows in fostering successful communication in his organization is in those interactions themselves. In the example of the workgroup, a productive meeting is the product of the group using the genre "the gathering" to help give the "right" shape and meaning to the interactions that take place in their weekly sessions: some of what they know in conducting productive meetings is in their interactions with one another.

Each of these is an example of dynamic affordance—of what becomes possible when knowledge is used as a tool in the context of situated activity. Each is also an example of the importance of both knowledge and knowing in understanding the role played by what we know in

organized human activity. It is by adding knowing to knowledge that we can begin to account for the relationship between what we know and what we do. And it is also how we can begin to see how new knowledge and knowing are generated.

Figure 4 shows the four forms of knowledge from Figure 3, the focus of the epistemology of possession, with a circle superimposed that represents knowing, the focus of the epistemology of practice. The arrows suggest active use of knowledge in our interaction with the social and physical world. Within this interaction lies what we have called the generative dance.

Knowing does not sit statically on top of knowledge. Quite the contrary, since knowing is an aspect of our interaction with the world, its relationship with knowledge is dynamic. Each of the forms of knowledge is brought into play by knowing when knowledge is used as a tool in interaction with the world. Knowledge, meanwhile, gives shape and discipline to knowing.⁸ It is this reciprocal interplay between knowledge and knowing that we call “bridging epistemologies.”

It is by bridging epistemologies that it is possible to draw among the four forms of knowledge within the same activity. Individual and group knowledge are both used, for example, in activities that dynamically afford both the practice of a given skill by an individual and “trying it out” by a group learning it—as when a choreographer teaches through demonstrations while a dance troupe follows. The group acquires tacit knowledge in practice as they develop a useful understanding, for example, of the moves employed in the piece through interacting with the demonstrations of the instructor (Cook 1982, Lave and Wenger 1991). It is within this interaction, moreover, that the troupe’s new knowledge (genres) and new forms of knowing (performing the dance) are generated (a generative dance—literally).


What we are proposing here is more than a shift in

language; it is a shift in focus from performing operations on existing knowledge to making something new. It is a shift in perspective that is meant to provoke different ways of assessing the role of what is known (both as knowledge and knowing) in an organization’s ability to learn, to maintain quality, to develop competencies, to innovate, etc. Organizations not only create knowledge, they also—and usually primarily—create goods and services. In doing so, they need to be increasingly innovative. And this requires, we believe, attention not only to what they possess, but also to how they practice. This calls for a broadening of focus from one epistemology to two, including the generative potential of interplay between them.

In this sense, the generative dance entails productive inquiry in a substantial and robust sense: it is not only productive as a team is productive when it meets a preset quota; it is truly *generative*. By this we mean that it is a source of innovation, of productive change—as when a team invents new ways of working more effectively. In a very basic sense, for example, the activity that conversation affords is not limited to a merely additive back and forth exchange of information. When Emma says to Andrew “I’ve been doing it this way,” Andrew not only adds that knowledge to his own, but he also takes it into the context of his own experiences, skills, sensitivities, and the like (and vice versa when Andrew makes his reply). By placing Emma’s knowledge into Andrew’s contexts, the conversation can evoke novel associations, connections, and hunches—it can generate new insights and new meaning. As everyone has experienced, a conversation’s back-and-forth not only dynamically affords the exchange of knowledge, it can also afford the generation of new knowledge, since each remark can yield new meaning as it is resituated in the evolving context of the conversation. Through conversation, Emma and Andrew can negotiate a joint understanding of what “doing it this way” means. This shared meaning, then, constitutes for them the genre “Emma’s way,” which, in turn, can become an innovative and more effective means to read, understand, and carry out their work together. In this way, conversation affords more than an exchange in which the net sum of knowledge remains the same; it dynamically affords a *generative dance* within which the creation of new knowledge and new ways of using knowledge is possible.

Engaging in such conversation is a practice that does epistemic work; it is a form of knowing. Knowing entails the use of knowledge as a tool in the interaction with the world. This interaction, in turn, is a bridging, a linking, of knowledge and knowing. And bridging epistemologies makes possible the generative dance, which is the source

Figure 4 Adding Knowing to Knowledge



of innovation. The generative dance, within the doing of work, constitutes the ability to generate new knowledge and new ways of using knowledge—which knowledge alone cannot do. And which the organizations of the future cannot afford to neglect.

4. Implications

We have found the perspective outlined above to have far-reaching implications for our work, in theory and in practice, and in assessing the work of others. Seeing each of the four forms of knowledge as unique, finding knowledge and knowing to be distinct, seeing how different epistemic work is done by different forms of knowledge and knowing, and understanding the notions of dynamic affordance and the generative dance—all this has not left our sense of how groups can and do work undisturbed. Below we briefly sketch out three cases that help make clearer some of the actionable and theoretically significant implications of this perspective.

The first case is drawn from Nonaka and Takeuchi's work on the "knowledge-creating company" (1995). Among their insightful explorations of "knowledge creating" is a case of a company's development of a bread-making machine. We build on their case, and argue that the perspective we have put forth here expands and makes more robust their notion of "knowledge creation." The second case deals with three Boston-area workshops that make world-class flutes. What the flutemakers know that enables them to make instruments of the highest quality, we argue, is found both in the knowledge they possess and in the ways they interact with the instruments and each other. The third case is a brief look at how a group of mechanical engineers in Xerox have created innovative *new* technologies in part through generative interactions with *old* mechanisms.

Machine Design

In their study of "the knowledge-creating company," Nonaka and Takeuchi (1995) illustrate what they call the "conversion" of tacit knowledge into explicit knowledge with the example of a company's development of a bread-making machine. A good bread-making machine must be able to knead dough properly. Yet, Nonaka and Takeuchi note, this is something "which is essentially tacit knowledge possessed by master bakers" (p. 63). So one of the company's software developers became an apprentice to a prominent hotel's head baker. She was then able, according to Nonaka and Takeuchi's interpretation, to "transfer" the tacit knowledge she acquired in working with the master baker to the engineers who were designing the machine's kneading mechanism by "converting" it into explicit knowledge "by using the phrase 'twisting

stretch'" (p. 104). The engineers used this knowledge in their work on the mechanism, and the software developer evaluated the results in a "trial-and-error process [that] continued for several months" (p. 104). Ultimately a good mechanism was produced. Nonaka and Takeuchi's argument, then, is that the tacit knowledge the software developer acquired by "observing and imitating the head baker" was converted into explicit knowledge through the use of the phrase "twisting stretch" (p. 105), which, along with the engineers' technological knowledge, enabled the group to produce a prototype of the machine (p. 106). In this way, they argue, the group was engaged in "knowledge creation."

We interpret this example somewhat differently. Yet, we believe an interpretation from the perspective of the generative dance serves to strengthen Nonaka and Takeuchi's central claims about "knowledge creation."

We see in the case the same distinct epistemological forms we saw in the bike-riding example, but now also at the organizational level. For us, the case is also an instance of bridging epistemologies, where the practices of the group (its ways of knowing) enabled it to draw simultaneously on different forms of knowledge possessed by different people. In this way, the individual tacit knowledge of the software developer and the explicit group knowledge of the engineers were both used by the team as a whole as tools within a productive inquiry (the trial-and-error process) that enabled them to design a successful kneading mechanism: various interactions by the group using specific tacit and explicit knowledge afforded the generation of both knowledge and new ways of knowing.

Following our interpretation, the example entails both "bridging epistemologies" and the "generative dance." In making the machine, the design team drew on all four types of knowledge (by bridging epistemologies). There was the explicit technical knowledge each member of the team possessed. We imagine that there were also explicit group stories or metaphors, since such are all but universally found in groups. Individual tacit knowledge comes into play in both the master baker's skill and in what the apprenticed developer acquired. And there was group tacit knowledge, we claim, in the form of the useful meaning that "twisting stretch" (as a genre) came to have for them (more on this in a moment).

In addition to the use of the different forms of knowledge, there was also knowing—that is, epistemic work that was part of the team's interaction with machine parts, bread dough, and each other. This interaction (this way of knowing) entailed use of the team's various bits of knowledge as tools. The interaction also involved dynamic affordance within which (alone) the team was able

to recognize and make use of the knowledge associated with the term “twisting stretch” (just as being able to ride a bicycle requires the dynamic affordance of being on a bicycle in motion in order to make use of the knowledge associated with “turn this way”). In particular, the term “twisting stretch” referred to both the individual tacit knowledge of the developer and the tacit knowledge of the group. Using the term in the trial-and-error process provided a way of going back and forth between the two. In essence, the term functioned as a kind of “boundary object” (Star and Griesemer 1989) that straddled bread making and machine making. Through the successive iterations of mechanism design, the engineers negotiated with the developer the proper meaning and use of the term in application to the motion from bread making that they were aiming to capture in a machine operation. In this way, the meaning of the term “twisting stretch” became a genre for the team as a whole (i.e., group tacit knowledge): it was the way they identified and understood the “right” movement in both bread making and machine making. By bridging knowledge and knowing in actual interaction with the machine and each other (that is, by treating knowledge as a tool of knowing), the team was able to use the term “twisting stretch” to draw on both individual and group tacit knowledge simultaneously in practice.

The generative dance can also be seen in the “twisting stretch” example. “Twisting stretch” as a genre (the shared meaning of the term), and the ability to use it in designing the prototype, were *new* things—a *new* bit of knowledge and a *new* way of knowing. They were not variant expressions of knowledge that already existed. They were created, we maintain, through the generative dance. That is, the design team used explicit and tacit knowledge as tools in interaction with machine parts and one another in an instance of productive inquiry that ultimately generated new knowledge and knowing. One of the team’s aims was for the engineers to acquire a sense of the proper kneading motion. This entailed interaction between the engineers’ machine making (a way of knowing) and the software developer’s tacit knowledge (associated with her bread making). This resulted in the *generation* of the genre “twisting stretch” (the group knowledge of what the term means). It was *not* tacit knowledge converted into explicit knowledge, it was *new* knowledge generated by the team. As a bit of knowledge, “twisting stretch” became a meaningfully useful tool in two forms of knowing: the software developer’s bread-making and the engineers’ machine making.

It is our focus on new knowledge and new knowing that leads us to prefer the concept of “generating” to that

of “converting” (as used by Nonaka and Takeuchi). “Conversion” tends to suggest an operation that is *applied* to knowledge rather than a concrete interaction with the world that generates knowledge. In converting feet to meters, an equation is applied to the measurement in feet and yields a measurement in meters, *without going back to the object at hand* to remeasure it. In our view, given one kind of knowledge, the only way to get the other is precisely *by going back to the object at hand and interacting* with it. For us, the “trial-and-error process” Nonaka and Takeuchi identify is an example of just this sort of interaction with the world. What the design team did was not a conversion process applied to the software developer’s tacit knowledge; it was an exercise in productive inquiry carried out by the group in interaction with bread dough, machine parts, and each other. This interaction dynamically afforded the use of both explicit and tacit knowledge, and ultimately generated new knowledge and a new way of knowing.

Flutemakers

The case of the three flute companies that manufacture world-class instruments allows us to take these notions further. They are particularly illustrative of the notion of dynamic affordance and its role in the generative dance.

The Boston workshops produce flutes that are embraced by the flute world as instruments of the finest quality. And the flutes of each workshop have a distinctive character recognizable by knowledgeable flutists as the flute’s “feel” (generally, how the instrument feels when it is being played—not, incidentally, how it sounds). Both the high standard of quality and the unique character of each brand of flute are highly valued by the flute world.⁹

For most of their history, each workshop has had between 20 and 40 flutemakers (including those who are owners and managers) plus one or two office staff. The flutemakers work in teams, each flute being the product of a number of flutemakers, with each flutemaker working only on part of the instrument. (It is rare that a single person has the ability at any one time to make an entire flute, although some work on numerous aspects of flute-making over the course of their careers.) A flutemaker, meanwhile, might work with a particular set of colleagues on one batch of flutes and with a different set on a later one. Over their history, the workshops have gone through generations of flutemakers (the oldest of the workshops dates from around 1900, the newest was established in 1977).

Because flutes are physical objects, the quality and character of each flute is inextricably tied to very fine degrees of dimension and tolerance in how their pieces work and fit together. Many of these dimensions and tolerances, however, are not known or used explicitly by the

flutemakers. Rather, they are set by judgments of hand or eye. Typically, each flutemaker works on his or her part of the flute until it meets his or her standard of appearance and/or feel. Then it is handed on to the next flutemaker, who judges the work of the first by his or her own standards. If the work is not “right,” it goes back to the previous flutemaker to be reworked until both are satisfied. Some measurement tools are used, such as calipers and feeler gauges; but even when a part is measured, it is also checked out by feel or by eye, which are the final courts of appeal.

When an apprentice joins a workshop there are many things he or she must learn (apprenticeships have taken up to five years). Elements of what needs to be learned reflect all four forms of knowledge. There are concepts and rules about the types of parts, how they are connected, which tools are used for which functions, and so on. There are the skills needed to make flutes with the “right feel.” These bits of explicit and tacit knowledge are learned and used by the individual apprentices just as they are used daily by master flutemakers.

At the group level, there are stories and metaphors used explicitly among flutemakers that help guide and coordinate their work. At one of the workshops flutemakers would argue that a piece of work or a new company policy ought to be “the way the old man would want it,” referring to the founder of the company (this continued long after “the old man” had retired and died). There are also genres that constitute the shared meaning of the “right way” to use certain equipment (feeler gauges, for example) or how to identify and understand what is wrong with a piece of work. When a part is handed back to a previous worker, for example, it can come with a comment such as “this is a clunky one.” The flutemakers then hand the piece back and forth discussing its “clunkiness.” This interaction with the piece and with each other dynamically affords a negotiation in practice as to what exactly “clunky” means in reference to the piece at hand and concerning what work needs to be done to it. When the meaning associated with “clunky” becomes commonly used by the flutemakers in recognizing, discussing and working on subsequent problems, it functions as a genre in that workshop.

The examples above reflect different forms of knowledge that fit the four categories of the traditional epistemology. But having such knowledge is only part of what is needed to make world-class flutes. Knowing is also required. Accordingly, it is typical for an apprentice to work on flutes starting on his or her first day in the shop: he or she engages in the practice of flutemaking, and begins to acquire not only knowledge but also ways of knowing. An apprentice may be told explicitly that “these

keys need to work more solidly.” But it is only through practice, through actual working jointly with other flutemakers on the piece, that he or she will “get a feel” for what “solidly” actually means in that shop (“solidly” could mean quite a different thing at one of the other workshops). When a master flutemaker says something such as “this is what we call clunky,” an apprentice can only know what that means by learning what it feels like—and a master flutemaker can only agree that an apprentice’s work ultimately feels right by feeling the piece.

This is also true of accomplished flutemakers: part of what they know is in the daily handing of pieces back and forth and negotiating that a piece of work looks or feels right. Interaction with the instruments and other flutemakers dynamically affords the use, in practice, of the different forms of knowledge possessed by the flutemakers, individually and as a group. Another part of what the flutemakers know, another part of their epistemic work, is in their interactions themselves. The genre “clunky” is a tool flutemakers use in their interactions with each other; it does the epistemic work of group tacit knowledge. Being able to recognize when “clunky” gives way to the “right feel” and being able to negotiate that with fellow flutemakers are also part of what flutemakers know, they are instances of epistemic work done as part of the practice of world-class flutemaking. And they are instances of knowing. The interaction with the instruments and among flutemakers also entails the generative dance; it is here that new knowledge and new ways of knowing are created. The back and forth between an apprentice and a master flutemaker, for example, dynamically affords two things at once: 1) the use, in practice, of existing tacit knowledge possessed by the master in judging the feel of the apprentice’s work; and 2) the *generation* of new tacit knowledge and new ways of knowing for the apprentice. This is an instance of the generative dance.

An apprentice acquires new tacit knowledge in his or her interaction with the instrument and with a master flutemaker, and those interactions also dynamically afford the master using his or her tacit knowledge as a part of the practice of flutemaking. That is, the apprentice’s *new* knowledge is *generated* in an interaction that has been given particular shape and form by the master’s use of his or her *existing* knowledge. While on the surface this can appear to be a *transfer* of knowledge from the master to the apprentice, we see it as an interaction with the social and physical world (flutemakers and instrument parts) in which the master’s knowledge is used and the apprentice’s knowledge is *generated*.

The importance of tacit knowledge and its dissemination in organizations are also topics emphasized by

Nonaka and Takeuchi (1995). For them this dissemination, including its role in the creation of new knowledge, occurs in a process they call “socialization.” They hold that “the sharing of tacit knowledge . . . is a limited form of knowledge creation” because unless tacit knowledge “becomes explicit, it cannot be easily leveraged by the organization as a whole.” They then contend that “Organizational knowledge creation is a continuous and dynamic interaction between tacit and explicit knowledge” (p. 70).

We propose three shifts that we believe build on and strengthen Nonaka and Takeuchi’s general insight. First, as we have noted in detail above, we contend that it is not possible, under any circumstances, for tacit knowledge to become explicit (or vice versa). We do hold, however, that one can be a useful tool in the generation of the other through productive inquiry.

Second, since we hold that explicit and tacit knowledge are generated and disseminated each in its own right, whether either can “be easily leveraged by the organization as a whole” depends, in our view, on the specific needs and resources that an organization has at hand in a given situation. The generation of explicit knowledge can, at times, be necessary to the dissemination of tacit knowledge (or even to making tacit knowledge more “easily leveraged by the organization as a whole”). However, this is determined by its usefulness as a tool in productive inquiry in a given situation, not by general characteristics of explicit and tacit knowledge, as Nonaka and Takeuchi suggest. If explicit knowledge is needed, then it is explicit knowledge that needs to be generated and made sharable; if tacit knowledge is needed, then it must be generated and made sharable (as we see in the flute case). Or both (as is found in the case of the bread-making machine).

Finally, for us, the production of new knowledge does not lie in “a continuous interaction between tacit and explicit knowledge” but rather in our interaction with the world. Specifically, it lies in the use of knowledge (explicit and/or tacit) as tools of productive inquiry (of the sort we have called “knowing”) as part of our dynamic interaction with the things of the social and physical world.

Paper Handling

The significance of interaction with the *physical* world to dynamic affordance and the generation of knowledge and knowing found particular meaning for us in a recent research project in Xerox. In this research, it was discovered that, for a group of design teams, interacting with *old* artifacts is often a source of insights that are valuable in designing *new* technologies.

As part of a broader research project, what is known in Xerox about the design of “paper paths” was examined.¹⁰ These are the various electromechanical devices that move blank paper from a paper tray through a copier, printer, fax machine, etc. as it is “marked” and then out of the machine as a printed page. These are surprisingly sophisticated devices, and there are often significant challenges in designing them as product cycles and technological innovations call for their evolution and change. This work is typically done by small teams composed mainly of mechanical engineers.

This expertise in paper path design is one of Xerox’s traditional core competencies. Yet, through the course of the recent research, we came to recognize how some very valuable aspects of this competency are also embodied in the paper path mechanisms themselves. With time, engineers can forget, retire, move on, and the like—including, over enough time, entire cohorts or generations of engineers. By one way of thinking, then, some features of a given paper path’s design and functioning, particularly subtle or sophisticated features, would no longer be available to Xerox. But the research revealed that when design teams sense that there is something in an old paper path that could be of use in designing new ones, they pull out the old one and begin to work with it. It is clear in this “working with” old mechanisms that the teams are after tacit knowledge, not explicit knowledge (they have the technical drawings for that). In fact, they refer to being interested in how the mechanisms “sound, feel, and work together” when in operation and when being assembled and disassembled.

This case complements Nonaka and Takeuchi’s bread-making machine example. In that example, what the engineers needed was *explicit* knowledge about the “twisting stretch” movement so they could design a mechanism that would replicate it. While in the paper path example, the engineers needed *tacit* knowledge about the feel, sound, and operation of older mechanisms, which they could use in designing new ones. Moreover, in the Xerox engineers’ interactions with the older mechanisms, tacit knowledge was leveraged by the organization as a whole without requiring the use of explicit knowledge.

This research has led us to believe that we need radically to rethink what is needed to create and support “core competencies.” Since part of Xerox’s paper path competency is embodied in old artifacts, design teams need to have the kind of “hands on” interaction with those artifacts that affords the recapture or (to follow our terminology) the regeneration of those particular bits of knowledge associated with that part of the competency. For the design team, this regeneration occurs as part of

group practice: their dynamic interaction with the old paper path apparatus affords the acquisition by the team of (tacit) knowledge about significant aspects of how the mechanism looks, feels, and sounds when it is operating well. It can also afford the identification of significant dimensions, tolerances, and functions (explicit knowledge) associated with the look, sound, and feel of proper operation.¹¹

We also believe there is a need to rethink how competency is distributed—in particular, how it can be found both in what individuals and groups know and in their practices. Part of Xerox's competency in paper handling is embodied in existing artifacts, part in knowledge people possess. Part also lies in the ability of design teams to interact with old artifacts in ways that afford the regeneration, for the team, of the knowledge associated with those mechanisms. That is, the ability of these groups to do this is also part of Xerox's paper handling competency.

A design team's practices also include the generation of knowledge new to the group. This can be seen, for example, in the case of genres. In the context of their interaction with old mechanisms, a team will identify (through negotiation in practice) which aspects of how a mechanism sounds, feels, and works are significant and which not. That is, bits of machine design and behavior will take on particular meaning (they will become genres), and those meanings will play a role in how the team frames, understands, or reads both their further interactions with the old mechanisms and their design work on the new one.

Finally, we would note that putting the knowledge associated with the older mechanisms in the context of new product design efforts results in more than adding old knowledge to new projects. It is a dynamic practice that can also afford the generation of new ideas and new ways of working—something that is not in the knowledge alone. Given this, we argue that understanding such things as the retrieval of "intellectual capital" solely as a matter of tapping into a knowledge base (that is, as solely concerned with knowledge) leaves untapped (as well as unsupported, unrecognized, and underutilized) the generative power of the practices associated with recapturing old knowledge.

5. Conclusion

This essay aims to broaden the existing understanding of what and how people know, as that relates to the epistemological dimension of organized human activity. We have offered the notions of distinct kinds of knowledge, productive inquiry, dynamic affordance, and the generative character of knowing to enrich such related themes

as organization knowledge, knowledge creation, knowledge-based organizations, the management of intellectual capital, knowledge work, etc. Clearly, the perspective we have proposed both suggests and would benefit from further theoretical and empirical work. Among the numerous areas where further work could be done are the following.

How might issues of core competency be broadened if we were to ask not only what knowledge is entailed, but also what forms of knowing (i.e., how particular groups use the knowledge they have or acquire)? We see the core competencies of the flute workshops, for example, to include, along with the four forms of knowledge distributed among individuals and groups, ways of knowing reflected in the interactions flutemakers have with each other and the instruments. Such knowledge and knowing are essential to the organizations' world-class status, yet they are also unique to each workshop, and therefore cannot be transferred from one company to another. (In fact, when accomplished flutemakers have moved from one workshop to another, they have had to undergo "retraining" in order to do work consistent with the new company's style and standards.) Thus, there is a need for a better understanding and better models of how this essentially non-transferable or "situated" dimension of knowledge and knowing, as elements of an organization's core competency, can be "generated in" (rather than "transferred to") other groups or organizations.

There is a need for more case studies of knowledge-creating organizations, knowledge work, and knowledge management that focus not only on the body of knowledge that an organization acquires, stores, and transfers. Equally important are the ways organizations can dynamically afford, within the situated practices of ordinary daily work, the productive inquiry essential to ongoing innovation.

There is also the very practical question of how training and educational programs can be redesigned. Such programs need to take as their aim both passing on knowledge to individuals *and* creating situations that help groups develop practices (ways of knowing) that make use of knowledge in new, innovative, and more productive ways.¹²

We hope that an expanded understanding of what and how people know can help provide an enriched, more robust way of assessing, supporting, and honoring the epistemological dimension of all "real work," which alone gives life and power to such concepts as core competency, knowledge creation, knowledge work, and intellectual capital.

Acknowledgments

For their careful reading of and valuable comments on earlier drafts of this work the authors are indebted to Johan de Kleer, Daniel Denison,

Paul Duguid, Larry Hickman, Kristian Kreiner, Charles F. Sabel, Edgar Schein, Sim Sitkin, Susan Stucky, Jeanne Vickers, Hendrik Wagenaar, Jay Zimmerman, and Betty Zucker. They are also indebted to the anonymous reviewers of this essay and, in particular, to Paul Adler for their exceptionally provocative and useful comments. Portions of the research that contributed to the writing of this paper were supported by a grant from the National Science Foundation (#9320927).

Endnotes

¹The term “epistemology” refers properly to the study of knowledge, including questions concerning what counts as knowledge and how bodies of knowledge can be systematically organized. More casually, it can also refer to knowledge and bodies of knowledge themselves (rather the way “ecology” can refer both to the study of environmental systems and to those systems themselves). We make use of both senses of the term (depending on the context).

²By “epistemic work” we refer to the work people must do to acquire, confirm, deploy, or modify what needs to be known in order for them to do what they do.

³We are indebted to Susan Stucky of the Institute for Research on Learning and to J.-C. Spender for the initial idea of this 2 × 2 table.

⁴Discussion of explicit and tacit knowledge has a long history and has not by any means come to consensus. The terms used, how they are related, and the realities they point to vary considerably. Ryle (1949), for example, cast the discussion in terms of what it means to “know *how*” and to “know *that*.” For some (including us) the two types of knowledge are seen as quite distinct, while others may see them as two ends of a continuum.

⁵The ontological status of groups has long been an unresolved issue. For our purposes, we take the view that not every action by a human collective can be meaningfully or usefully reduced to an account of actions taken by the individuals in them (as the practices of coaches, orchestra conductors, and organizational managers would suggest). To this extent, we believe collectives can be coherently and usefully considered in their own right with respect to actions they perform and with regard to the possession of any knowledge used in those actions.

⁶Schon (1983), whose work also draws strongly on Dewey, makes a similar distinction in discussing what he sees as the need to shift from pure technical rationality to what he calls “reflection-in-action” in professional practice.

⁷This sense of significant properties arising in the interaction between the self (or group) and the world is also a central theme in the work of the twentieth-century Japanese philosopher Watsuji (1961).

⁸Our language here (and at other points) suggests a resonance with structuration theory, especially with Giddens (see, for example, Giddens 1979, especially chapter 2; and Cohen 1989, especially chapter 1). Structuration theory’s treatment of praxis as constitutive of social structure, while social structure informs praxis, parallels our characterization of knowledge as brought into play by knowing, while knowing is disciplined by knowledge. Some might reject any such parallel, given that our focus is essentially epistemological, while structuration theory (particularly Giddens himself) deliberately eschews epistemological concerns in favor of ontological ones. Others may see our treatment of the interaction of knowledge and knowing as an instance of structuration. For our part, we find the parallel a provocative one, both epistemologically and ontologically. Although a systematic consideration of this similarity is not within the scope of this essay, we

would make the following observations. We do not take the relationship between knowledge and knowing to be nothing more than a straightforward example of the more general relationship between structure and agency found in structuration theory (if for no other reason than that we believe neither structuration theory nor pragmatism makes the other epistemologically and/or ontologically redundant). At the same time, we believe that a fuller investigation of pragmatist epistemology and structuration ontology could find in the practice of productive inquiry a way to help the epistemological more fully rejoin the ontological within the purview of structuration theory.

⁹A fuller presentation of this case focusing on organizational learning can be found in Cook and Yanow 1993. An extensive presentation and analysis of the case, focusing on tacit skills, judgment, and apprenticeship within the cultural context of groups can be found in Cook 1982.

¹⁰This research was conducted as part of a project headed by Robert S. Bauer of Xerox Corporation and Estee Solomon Gray of Congruity. We are indebted to them for this example and for the project’s influence on our thinking in general.

¹¹In addition to innovation, the use of older artifacts can also be seen in the case of training. Clark and Wheelwright (1992) have observed that Braun maintains a collection of their old products for use in training new product designers.

¹²The theories and practices of “progressive education” might offer some provocative points of reference in this regard.

References

- Anderson, R., W. Sharrock. 1993. Can organizations afford knowledge? Xerox Technical Report EPC-92-104.
- Argyris, C., D. A. Schon. 1978. *Organizational Learning*. Addison-Wesley, Reading, MA.
- Bourdieu, P. 1977. *Outline of a Theory of Practice*. Cambridge University Press, Cambridge.
- Brown, J. S., A. Collins, P. Duguid. 1989. Situated cognition and the culture of learning. *Educational Researcher* 18(01) 32–41.
- , P. Duguid. 1991. Organizational learning and communities-of-practice: Toward a unified view of working, learning, and innovation. *Organ. Sci.* 2 40–57.
- , ———. 1994. Borderline issues: Social and material aspects of design. *Human-Computer Interaction* 9.
- Clark, K. B., S. C. Wheelwright. 1992. *Managing New Product and Process Development*. Free Press, New York.
- Cohen, I. J. 1989. *Structuration Theory: Anthony Giddens and the Constitution of Social Life*. St Martin’s Press, New York.
- Cohen, M. D., P. Bacdayan. 1994. Organizational routines are stored as procedural memory: Evidence from a laboratory study. *Organ. Sci.* 5(4) 554–568.
- Cook, S. D. N. 1982. Part of what a judgment is. Ph.D. dissertation. Massachusetts Institute of Technology (available from the author or from University Microfilms).
- . 1994. Autonomy, interdependence and moral governance: Pluralism in a rocking boat. *Amer. Behavioral Sci.* 38(1).
- , D. Yanow. 1993. Culture and organizational learning. *J. Management Inquiry* 2(4).
- Gaver, W. W. 1991. Technology affordances. *Proceedings of CHI ’91*. ACM Press, New Orleans. 79–84.
- . 1996. Affordances for interaction: The social is material for design. *Ecological Psych.* 8(2) 111–129.

- Gibson, J. 1979. *The Ecological Approach to Visual Perception*. Houghton Mifflin, New York.
- Giddens, A. 1979. *Central Problems in Social Theory: Action, Structure and Contradiction in Social Analysis*. University of California Press, Berkeley and Los Angeles.
- Hamel, G., C. K. Prahalad. 1994. *Competing for the Future*. Harvard Business School Press, Boston.
- Hickman, L. 1990. *John Dewey's Pragmatic Technology*. Indiana University Press, Bloomington.
- Hutchins, E. 1991. The social organization of distributed cognition. L. B. Resnick, J. M. Levine, S. D. Teasley eds. *Perspectives on Socially Shared Cognition*. American Psychological Association, Washington, DC.
- Kogut, B., U. Zander. 1996. What firms do?: Coordination, identity, and learning. *Organ. Sci.* 7(5).
- Krogh, G. von, J. Ross. 1995. *Organizational Epistemology*. St. Martin's Press, New York.
- Lave, J., E. Wenger. 1991. *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press, Cambridge.
- Leonard-Barton, D. 1995. *Wellsprings of Knowledge: Building and Sustaining the Sources of Innovation*. Harvard Business School Press, Boston.
- March, J. G., J. P. Olsen. 1976. Organizational learning and the ambiguity of the past. *Ambiguity and Choice in Organizations*. Universitetsforlaget, Oslo, Norway.
- Nonaka, I. 1994. A dynamic theory of organizational knowledge creation. *Organ. Sci.* 5(1).
- , H. Takeuchi. 1995. *The Knowledge Creating Company*. Oxford University Press, New York.
- Orlikowski, W. J., J. Yates. 1994. Genre repertoire: The structuring of communicative practices in organizations. *Admin. Sci. Quart.* 39.
- Orr, J. E. 1996. *Talking About Machines: An Ethnography of a Modern Job*. Cornell University Press, Ithaca.
- 1990. Sharing knowledge, celebrating identity: War stories and community memory in a service culture. D. S. Middleton, D. Edwards, eds. *Collective Remembering: Memory in Society*. Sage Publications Limited, London.
- Ortega y Gasset, Jose 1961a [orig. 1941]. History as a system. *History as a System*. W. W. Norton, New York.
- 1961b [orig. 1941]. Man the technician. *History as a System*. W. W. Norton, New York.
- Polanyi, M. 1983 [orig. 1966]. *The Tacit Dimension*. Peter Smith, Magnolia, MA.
- Rorty, R. 1982. *Consequences of Pragmatism*. University of Minnesota Press, Minneapolis.
- Ryle, G. 1949. *The Concept of Mind*. Hutcheson, London.
- Schon, D. A. 1983. *The Reflective Practitioner: How Professionals Think in Action*. Basic Books, New York.
- Simon, H. A. 1991. Bounded rationality and organizational learning. *Organ. Sci.* 2(1).
- Sims H. P., Jr., D. A. Gioia, and associates. 1986. *The Thinking Organization; Dynamics of Organizational Social Cognition*. Jossey-Bass, San Francisco.
- Sitkin, S. B. 1992. Learning through failure: The strategy of small losses. *Res. Organ. Behavior* 14.
- Spender, J.-C. 1996. Competitive advantage from tacit knowledge?: Unpacking the concept and its strategic implications. B. Moingeon, A. Edmondson, eds. *Organisational Learning and Competitive Advantage*. Sage, London.
- Star, S. L., J. R. Griesemer. 1989. Institutional ecology, 'translations' and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39. *Soc. Stud. Sci.* 19.
- Stewart, T. A. 1997. *Intellectual Capital: The New Wealth of Organizations*. Doubleday/Currency, New York.
- Turner, S. 1994. *The Social Theory of Practices: Tradition, Tacit Knowledge, and Presuppositions*. University of Chicago Press, Chicago.
- Vickers, G. 1976. Technology and culture. Invited paper given at the Division for Study and Research in Education, Massachusetts Institute of Technology, Cambridge, MA (MS available through S.D.N. Cook).
- Watsuji, T. 1961. *Climate and Culture*. Monbusho, Tokyo.
- Weick, K. E. 1991. The nontraditional quality of organizational learning. *Organ. Sci.* 2(1).
- , F. Westley. 1996. Organizational learning: Affirming and oxymoron. *Handbook of Organization Studies*. SAGE Publications, Thousand Oaks.
- , Roberts, K. H. 1993. Collective mind in organizations: Heedful interrelating on flight decks. *Admin. Sci. Quart.*, 38 357–381.
- Wenger, E. 1997. *Communities of Practice: Learning, Meaning and Identity*. Cambridge University Press, Cambridge.

Accepted by Paul S. Adler; received December 12, 1996. This paper has been with the authors for one revision.