

# Chapter 9

## Learning on the Shop Floor: The Behavioural Roots of Organisational Knowledge

John D. Hanson

**Abstract** If it is accepted that knowledge is a valuable asset of a firm, then it follows that learning or the ability to learn is equally important. This study used shop-floor observations to understand how this learning takes place in an organisational setting. It was discovered that for improvements in organisational knowledge to be realised, two distinctly different components had to be present. The first component consists of the performative aspects of the task at hand: techniques, resources, scientific facts, and the like. These elements are readily codified, stored, retrieved, and copied. They are cumulative in the sense that acquiring a new technique does not eliminate the ability to use an existing appropriate one. The second component is the underlying logic by which techniques are selected and applied to achieve a goal. This component corresponds closely to Plato's description of knowledge as justified belief. This component is not readily observable and is not cumulative in the sense that one cannot simultaneously hold conflicting beliefs. Because the context and justification of belief are defined by cultural norms, this introduces a strong behavioural component to organisational knowledge. The existence and importance of this largely invisible component have significant implications for managers who wish to promote organisational learning and for researchers who wish to study it.

---

J. D. Hanson (✉)  
School of Business Administration, University of San Diego, 5998 Alcalá Park,  
San Diego, CA 92110, USA  
e-mail: hansonj@sandiego.edu

## 1 Introduction

Few would challenge the premise that knowledge is an important asset of a firm. There is also a corresponding belief that this importance to the firm and to the economy at large is increasing with time, attributable primarily to the increasing technical complexity of our world (Drucker 1993; Bettis and Hitt 1995; Mukherjee et al. 1998). That being the case, it must also be true that learning (in the general sense of improving the state of knowledge) is a process critical to firm success.

The recognition of knowledge as an important *resource* of the firm is not new (it can be seen clearly in Penrose 1959 for example), but increased attention to this fact has coincided with the emergence of the *Resource-Based View* of the firm (RBV), a term introduced by Wernerfelt (1984), in the field of strategy. Arguing the economic origins of the RBV, Williamson (1981) makes the case that knowledge is a source of economic rent. Extending this reasoning, Grant (1996) proposed the Knowledge-Based View (KBV) of the firm. Central to this view is the argument that knowledge is actually a firm's only source of economic rent on the basis that it is the only resource whose use incurs no opportunity cost. Many researchers have made statements to this effect (Prahalad and Hamel 1990; Barney 1991; Nelson 1991; Leonard-Barton 1992; Cyert et al. 1993; Henderson and Cockburn 1994; Nonaka 1994; Bates and Flynn 1995; Nonaka and Takeuchi 1995; Kogut and Zander 1996; Miller 1996; Spender 1996; Davenport and Prusak 1998; Nahapiet and Ghoshal 1998).

One result of this level of attention is a vast literature on organisational knowledge in which the subject is approached from a variety of viewpoints with an equally wide range of findings. To clarify the discussion, this article will adopt a narrow definition of the term organisational knowledge that addresses both parts of the term: what is knowledge, and what is it that makes it organisational in scope (as opposed to common or individual knowledge)?

A key issue in knowledge is the *appropriability* of the rents that derive from it (Collis and Montgomery 1998). In that context, it is an important distinction whether knowledge is uniquely a property of the organisation (which would then capture the rents), or a property of the individuals who happen to make up the organisation at the time (and who should therefore be able to capture all or most of those rents through above-normal salaries). In the work that follows, the term *organisational knowledge* will be used to mean a form of knowledge that is uniquely a property of the firm; something that the firm can claim as its own and is robust against disturbances such as employee turnover.

The question of what constitutes knowledge is more contentious. Nonaka and Takeuchi (1995) note that modern Western philosophy has been dominated by the *Cartesian Split*, the paradigm of Descartes in which knowledge, and by extension truth, are independent absolutes which can be discovered and are therefore independent of the person. Polanyi (1958) however concluded that this view is ultimately not sustainable and that knowledge is, in the final analysis, something unique to the knower. In this he falls back on a definition of knowledge as *justified*

*belief*, a definition that Nonaka and Takeuchi (1995) attribute to Plato. It is the role of belief and the context of justification that makes this an individual or personal form of knowledge.

However, our interest is primarily in knowledge as a property of an organisation, and it is not clear that definitions centred on individuals are relevant when applied to organisations. Cyert and March (1963) first proposed that an organisation could learn independently from its individual members but did not leave us with a comprehensive definition of organisational knowledge. Grant (1996), in proposing the Knowledge-Based View of the firm, avoided the issue by ascribing knowledge purely to the individuals comprising the organisation. Teece et al. (1997) also circumvented the issue by arguing that what the organisation possesses is the complementarity of the knowledge of individuals assembled within the organisation. Nonetheless, the term *Organisational Knowledge* is very much part of our vocabulary and many feel that it is a distinct form of knowledge. Specifically, Hedberg (1981) describes it as being a “robust property” of an organisation.

To arrive at a useful definition of organisational knowledge, we follow the lead of Pentland (1992), who drew on the philosophy of American pragmatism, and specifically the work of George Herbert Mead, William James and John Dewey. This school also rejects the Cartesian view of knowledge and it is from their writings (specifically: Dewey 1916; Dewey and Bentley 1949) that we derive the idea that knowledge is *situated*, that is to say: defined only relative to a situation. Accordingly, we adopt the definition of organisational knowledge as *situated performance*, which means the ability to deliver a satisfactory performance in a defined set of circumstances. This can be justified by the fact that organisations generally have a defined purpose that defines the situation.

Many similar definitions or statements about organisational knowledge have been offered. Amin and Cohendet (2000) speak of the firm as: *a social institution, the main characteristic of which is to know well how to do certain things* (p. 96). In the same vein, Orlikowski (2002) equates knowledge with *effective action* as do Sabherwal and Becerra-Fernandez (2003).

With knowledge defined as a form of capability, there is still the issue of how it becomes a property of the firm; where it resides, as it were. Here we turn our attention to the subject of *routines*. We do so because many researchers have concluded that routines are the means by which knowledge is embedded in an organisation (Levitt and March 1988; Cohen 1991; Amundson 1998; Orlikowski 2002). While no single researcher can claim priority on the concept of routines as organisational knowledge, the idea is clearly present in the work of March and Simon (1958) who observed that what we are calling organisational knowledge existed in the form of *performance programs*, which we now more commonly refer to as *routines*. These routines are fundamental to production and operations management. Although they drew our attention to routines, March and Simon (1958) did not describe how they were encoded in organisations or how they were created or changed.

Progress has been made on the anatomy of routines by Feldman and Pentland (2003) who expanded the basic concept of routines to include the potential for

variation and selection, and therefore evolutionary, incremental change. More recently, Feldman and Pentland (2005) and Pentland and Feldman (2005) have argued that to be understood, routines must be considered as consisting of two layers; the *ostensive* and the *performative* (plus associated *artefacts*). The performative layer is the observable sequence of actions that constitute the routine as it is practiced in the moment. To an outside observer, it might seem that this would be a complete description of the routine. The ostensive part is the relatively unobservable underlying purpose or intent of the routine. The need to understand the ostensive layer was motivated by their observation that the performative layer tended to vary unpredictably (to our outside observer at least) in response to changing conditions. If the routine is the embodiment of organisational knowledge, the ostensive part is at least as important a determinant of situated performance as the performative part. It is the feedback loop by which the performance is assessed relative to the situation and by which the process is adjusted to suit.

By extension, if performative and ostensive components are both essential components of organisational knowledge, then beyond some point, organisational learning must involve changes to both parts. This is precisely the point made by Argyris and Schön (1978) when they coined the term “*double loop learning*” to describe the case where the ostensive part is examined and reconsidered. They did not use the term ostensive; instead they spoke of the “theory in use” that provided the belief that a certain set of actions would result in the desired outcome.

In Operations Management, the literature has focused heavily on performative aspects (“tools”) of routines, with a resulting lack of explanatory power. To cite a well-regarded paper as an example, Flynn et al. (1995) analyse the performative aspects of organisations that can be discerned from survey data, and obtain results that hint strongly at the importance of an (unidentified) underlying ostensive aspect. Using different terminology, Miller (1996) showed recognition of this gap by stating that competitive advantage lies: “not in specific resources or skills, but in orchestrating themes” (p. 509).

With organisational knowledge defined as situated performance we can derive an operational definition of *organisational learning* as: *an increase in situated performance*. We should probably focus on deliberate rather than accidental improvements in performance, so this definition can be supplemented by one proposed by Argyris and Schön (1978): *the detection and correction of error*. Equivalence of the definitions can be established by regarding “error” to be any unsatisfactory level of situated performance. Significant in this definition, and very much to the authors’ central point is that the assessment of “error” depends not only on the level of performance, but also on the perception of what it should be and how it should be measured.

Whether this action of learning consists of creating new knowledge or applying existing knowledge may be an artificial distinction (Starbuck 1992) but what we can know is that the end result is an improved capability on the part of the firm. However, while the knowledge created may be a property of the firm, the actions taken are by individuals who must be presumed to respond to behavioural

influences and limitations. If we are to understand the mechanism by which knowledge gets into practice, we need to understand more about why people do what they do.

## 2 The Shop-Floor Perspective

Researchers are generally obliged to study firms from the outside. As a result, it is not surprising that much of what we know about organisational learning is expressed in terms that describe the behavior of the firm as a monolithic entity. An important example of this is the term *Absorptive Capacity* (Cohen and Levinthal 1990). This is a measure of a firm's ability (or perhaps: propensity) to learn, and while it may provide a useful way to categorise and describe firms, it tells us relatively little about how this takes place. When a firm is seen as a single entity, the analogies to an individual person are inviting and Cohen (1991) and Cohen and Bacdayan (1994) make the case that, just as individuals store their learning in either procedural or declarative memory, so too do organisations have distinctly different modes of retaining their knowledge. Organisational routines are argued to be the analogue of procedural memory—the largely unconscious “how to” that functions relatively independent of logic or data. Again, while this provides useful labels, the terms tell us little about how the knowledge got there in the first place or how it is changed.

When we attempt to look inside the organisation, typically through the use of survey instruments or structured interviews, we encounter an additional difficulty in that such data is almost always collected at a single level within an organisation. In a recent study on innovation, Rothaermel and Hess (2007) noted that the antecedents to innovation took place at the individual, firm, and network levels. To focus on only one level, they noted, could lead to two serious problems, or as they put it:

First, concentrating on only one level of analysis implicitly assumes that most of the heterogeneity is located at the chosen level, whereas alternative levels of analysis are considered to be more or less homogeneous. Studies of firm-level heterogeneity assume, for example, that significant variation occurs at the firm level of analysis, whereas individuals are more or less homogeneous or randomly distributed across firms. Second, when focusing on one level of analysis, researchers implicitly assume that the focal level of analysis is more or less independent from interactions from other lower- or higher-order levels of analysis. Firm-level heterogeneity, for example, is assumed to be relatively independent from individual- or network-level effects. *Taken together, the assumptions of homogeneity in, and independence from, alternate levels of analysis are serious concerns that could lead to spurious empirical findings* (p. 899, italics in the original).

Given these limitations on existing research, it becomes apparent that to gain additional insight into the process of organisational learning it will be necessary to adopt different approaches. In this we have some guidance from Ohno (1988) who was adamant that learning takes place where the action is, hence *genchi genbutsu*

(roughly translated as: “get out on the shop floor and see what is really going on,” (Dennis 2002, p. 141).

Doing so introduces an ethnographic component to the research, in which the output is *cultural interpretation* as opposed to simple data (Hoey 2011). The nature and goal of such research is to gain enhanced insight, at the expense of wider generalizability. The introduction of a cultural dimension raises issues of methodology and interpretation, summarised by the *etic/emic* divide (Martin 2002, p. 36). The distinction in this divide is that the researcher’s point of view is either within the culture, or outside of it. Each point of view will fail to see things that the other will. Most typically, the etic (outsider) view will fail to see the significance attached to various symbols or actions, while the emic (insider) view will fail to see that certain taken-for-granted truths are not universal. In addressing the shortcomings of the etic view, Bartel and Garud (2003) note that: “to see and understand narratives, the researcher must become *semi-native*” (p. 337, emphasis in the original). This argues for the selection of a participant-observer model of observation (Martin 2002, pp. 48, 210) which can bridge the “insider” and “outsider” views of an organisation.

### 3 The Study

The study described here was motivated by the objective of shedding light on the process of organisational learning and sought to do so observing the process in action. This required a suitable setting and an appropriate method. As discussed above, the participant-observer method was felt to be appropriate for this question because the goal was not simply to record actions or results, but to understand the logic behind the actions, as interpreted by the subjects themselves.

A setting had to be chosen in which this model could be employed effectively. For efficiency, it was necessary to select settings in which organisational learning was expected to occur. For validity it was necessary to narrow those down to ones that would be relatively free of confounding influences. An additional level of validity is established by repetition: there are numerous examples of in-depth research of this kind that result in single-sample anecdotes, whose relevance to each other is hard to establish. Ideally what was wanted was a group of settings that conformed to some consistent standards so that we could compare the actions of different (but comparable) people in different (but comparable) organisations addressing different (but comparable) problems.

The class of activities known as Kaizen Events, as described by Melnyk et al. (1998), was selected as satisfying these requirements. These events have useful characteristics for the purposes of this study. By their very nature they are designed to promote and capture organisational learning. They are also bounded in their scope, their duration and their slate of participants. This bounding removes many unobservable factors that could otherwise confound the observations. A key choice was whether to seek out Kaizen events from different corporations or to

study events taking place under a single corporate umbrella. The latter alternative was eventually chosen to take advantage of the greater uniformity of approach. The common mode problem was minimised as much as possible by insisting that the events studied should take place in distinctly different operating units of the corporation.

A target corporation (hereafter: the Corporation) that was known to conduct significant numbers of Kaizen events was approached and asked for permission to observe some of the events. The Corporation, through its individual operating companies, is engaged in the manufacturing of a wide range of products used in commercial and residential construction and remodelling. At the time of the study, there were over 50 of these autonomous operating companies, organised into five product-based business lines.

It was agreed that a researcher would be allowed to attend a number of events acting as a participant-observer. Within the Corporation, there was a range of events to choose from, but many were not totally suitable for the study. Only events focused on process improvement in manufacturing operations were considered, and any event that was conducted in two parts or that represented a continuation of an earlier event was ruled out. This two-part structure was fairly widely used in the Corporation, and while it has certain practical advantages, it compromised the ideal of a bounded event because the actions of the participants could not be observed in the intervening periods. As a result of this selection process, three research cases were identified. The Kaizen events comprising these cases took place in a six-month time window, during which there were no major shifts in the economic environment experienced by these companies. The key features of these three events are summarised in Table 1.

There are obvious limitations to this approach. Organisational learning is only being observed in one particular type of setting and it would be improper to assert that the observed mechanisms would broadly applicable to other situations. The objective however, was insight, not generalizability and method of data collection and analysis was designed to support this. The collection and presentation of data presents some difficulties in this type of research because the approach is *inductive*. By that, it is meant that we observe evidence (outcomes) that would verify potential hypotheses in order to arrive at more general propositions.

The issue is with the method by which observations (representing potential hypotheses) are distilled into propositions. Hunt (1991) offers two choices, neither of them entirely satisfactory. If the observations are quantifiable, and enough data can be collected, then propositions can be derived statistically, as in exploratory factor analysis. This is a somewhat weak method, since propositions derived in this way cannot be falsified; one can only say that they are supported or not within certain confidence limits.

When observations are not quantifiable or are few in number, both of which are the case in this research, the remaining alternative is pattern matching, which is the search for distinctive and recurring patterns in the data (Hunt 1991). The realities of the social sciences dictate frequent use of the method, and it is advocated in this context by Kaplan (1964). While this method can generate useful insight,

**Table 1** Summary of events

Description of process	Central goal of event	Team composition	Summary of major results	Comments
Event #1—Plumbing valve body fabrication, automated brazing carousel	Reduction of setup time to change fixtures for part # changes	<ol style="list-style-type: none"> <li>1. Internal facilitator (manufacturing engineer)</li> <li>2. Outside facilitator</li> <li>3. Tooling engineer</li> <li>4. Assembly worker</li> <li>5. Lead hand/setup specialist</li> <li>6. Warehouse supervisor</li> <li>7. Tool and die maker</li> <li>8. Quality control rep.</li> <li>9. Researcher</li> </ol>	<p>Labour content to change fixtures reduced by ~50 %. Numerous other small improvements, most significant: improved post-cooling to reduce sticking on fixtures, improved airflow (cooling) to reduce breakage of brazing wire during automatic feeding</p>	<p>Not all opportunities to reduce machine downtime were implemented. Time reduction was not validated in practice but is believed to be accurate. No scheduling changes were implemented at this time</p>
Event #2—Metal fabrication facility; cutting, bending, machining, welding, painting of machine frames	Reduce work-in-process inventory between welding and painting operations, also reduce floor space	<ol style="list-style-type: none"> <li>1. Director of operations (internal facilitator)</li> <li>2. Outside facilitator</li> <li>3. Quality control rep.</li> <li>4. Operator: assembly area</li> <li>5. Operator: cable area</li> <li>6. Operator from paint line</li> <li>7. Researcher</li> </ol>	<p>Developed revised painting schedule and “supermarket” rack between operations to reduce WIP by 75 % by cycling through paint colours four times/shift instead of once. Minor other housekeeping improvements</p>	<p>WIP and floor space reduction dependent on scheduling, not necessarily sustainable. Note that director title may mislead; manufacturing engineer would have been more typical designation</p>

(continued)



**Table 1** (continued)

Description of process	Central goal of event	Team composition	Summary of major results	Comments
Event #3—Final assembly and packing of shower control valves (retail and contractor packs)	Replacement of carton packing bench to be more ergonomically sound. Rearrangement of component storage to facilitate SKU changes	<ol style="list-style-type: none"> <li>1. Internal co-facilitator from human resources</li> <li>2. Internal co-facilitator; production supervisor</li> <li>3. Production team leader for area, first shift</li> <li>4. Production team leader for area, second shift</li> <li>5. Chem lab tech., 3rd shift</li> <li>6. Test technician</li> <li>7. Assembly operator</li> <li>8. Researcher</li> </ol>	<p>Revised bench designed, built and implemented. Stock of components rearranged by family to reduce handling and distance travelled during changeovers. Expected to improve housekeeping, reduce errors and damage (not quantified). Mistake-proofing implemented to prevent recurrence of a known issue</p>	<p>Key difference from other events is lack of quantitative targets—seemed to reflect higher level of confidence that team would do the right thing</p>

Hunt (1991) points out that it suffers from a lack of intersubjective confirmability. By this he means that if two researchers extract different patterns from the same data, there is no test to determine which is better. Under these circumstances, the best that one can hope for is a degree of consensus. In this research this issue was addressed through the use of a panel of experts who were asked to comment on the validity of some of the key patterns observed.

Since validity is of critical importance in case study research, particularly internal validity (Anderson et al. 1999) additional steps were taken. Purposeful case selection has already been described, as has *validation* by the expert panel. As recommended by Yin (1994), *triangulation* was employed in the form of the use of multiple sources of evidence to address the research question. This was possible in these cases because, in addition to the observations of activities, the full range of data and documents used by the project teams was made available to the researcher. To guide the actual collection of observations a ten-item *research protocol* (Ellram 1996), was developed from a comprehensive review of the literature on organisational learning. This was used as an event checklist to ensure that each case was addressed systematically in search of evidence supporting or contradicting propositions derivable from the literature. The actual transcripts taken during the events were augmented with additional details filled in at the end of each work day. Copies of these notes were returned to the team leaders after the event for *verification* and correction of any errors of fact. Once approved, axial or pattern coding was used (Miles and Huberman 1994). That is, each item of each transcript was reviewed to determine whether it constituted evidence for the presence or absence of any of the protocol elements. These observations were then sorted to provide a summary of the relevance of the protocol element in explaining what had taken place.

## 4 Key Findings: The Nature of Organisational Learning

As described above, each of the case settings involved a conscious effort to effect process improvements which, if successful, would be evidence that organisational learning had taken place (in the sense that *situated performance* would have been improved). Across the three settings, 38 identifiable process improvements were implemented, but our interest is not in the improvements themselves, rather the question is how they originated and came to be implemented.

It was initially supposed that this would involve examples of knowledge transfer, given that the central theme was the implementation of Lean techniques that are well-documented. In that view, organisational learning would consist of copying knowledge from one “bin” to another (Walsh 1995) or converting from one type to another (Nonaka 1994). As a result, part of the data collection was a tabulation of where people turned for inspiration when faced with a problem.

These results are summarised in Table 2, but this is incomplete because in every instance the first place everyone turned was to their own experience to see if

**Table 2** Source of information

Case #	Problem-solving examples	Consult with experts	Observe existing examples	Look up documentary records
1	17	4	3	1
2	14	3	1	1
3	7	0	1	0
Total	38	7	5	2

Experts are individuals outside the project team with experience relevant to the problem at hand  
 Existing examples are previously—implemented problem solutions elsewhere in the company  
 Documentary records are company records specific to the process or product (as opposed to supplier catalogs, for example)

they had seen something like this before that could be applied in the present situation. Since this was universal, it is not listed in the table. The second choice was to tap into the experience of others by seeking out experts with relevant experience. The key point in this table is that these first two categories (own experience and experiences of others) are *efficient* in the sense that they contain not only specific details, but context and interpretation as well. The last two categories are *inefficient* because they lack these things. In fact, documentary sources, even when apparently directly relevant to the problem at hand, proved to be generally unusable unless an expert could be found who could explain what was done, why, and the story behind the story. As a result, use of these resources was avoided even when they were available. As a simple example of this, in the first case a question arose as to whether a certain modification to a gas line would be allowed by the code. Now the text of the code would have provided a definitive answer, but the response was simply to call up a plumbing contractor that the company used regularly and ask him the question. He was able to answer this immediately, and more importantly he was also able to answer the real question which was whether or not an inspector would pass it, thereby addressing issues of context and interpretation.

The use of existing examples is an interesting situation, and an example will serve to illustrate the mechanism involved. In the third case, one of the tasks was to replace a packing bench with a new one that would be ergonomically friendlier. (The old one required operators to walk around it repeatedly and to reach excessive distances). It was observed that a bench on another assembly line in the plant had a good reputation in this regard, and two team members were dispatched to study it with the goal of replicating its key features. Since they could not copy it exactly, what they had to do was essentially to reverse-engineer it to discover the design rules for creating an ergonomically friendly work station. This they did in a very tacit way, taking only a few measurements and writing nothing down. Whether they discovered all of the logic embedded in the original design is unknowable and it is reasonable to assume that they also had a few preconceived ideas of their own about how such a bench should be built. When they were finished, the new bench was an improvement, so we can say that learning took

place. However, it would not be fair to call this a transfer of knowledge; more realistically we would have to describe it as a creation of new knowledge, albeit heavily influenced by existing example.

This example serves to illustrate the central theme of the findings: that knowledge is not transferred or converted; it is created anew in a form unique to its situation and from a combination of factors. This is fully consistent with the definition of organisational knowledge as situated performance, and is also analogous to the concept of situated cognition (Rosch and Lloyd 1978). Furthermore, each instance of organisational knowledge consists of two distinct parts. This insight is also reflected in the literature, but most explicitly by Pentland and Feldman (2005). In describing routines as the unit of analysis within which organisational knowledge is embedded, they characterised routines as being comprised of tangible, observable elements such as work practices and procedures (in this case, the bench itself), and an unobservable component they called the ostensive element (the design rules, and if necessary, the operators' understanding of how to use it). To be more precise, they characterised routines as containing three parts, the third being the physical artefacts such as fixtures or line layouts. For our purposes these serve much the same purpose as the performative elements such as work instructions. The ostensive component is hard to define precisely, but corresponds to the conceptualisation of the problem to be solved and the appropriate range of approaches for doing so. What was found, as illustrated in the above example, is that both parts had distinct impacts on situated performance. We can state this more broadly as a proposition:

*P1: Organisational knowledge (and by extension, organisational learning) must be understood as consisting of two distinct components, the performative and the ostensive.*

Note that this is similar to the often stated distinction between knowing "what" and knowing "how", that Edmondson et al. (2003) also characterise as a distinction between codified and tacit knowledge. In terms of their impact on performance, the terms performative and ostensive are preferred here because they are more descriptive of what is actually going on.

An example was presented above showing how both elements were essential to the creation of an effective work station. Armed with this insight, every instance of changes to routines (learning) could be analyzed the same way and in many cases the linkage between ostensive and performative elements was obvious. It is the exceptions that are interesting however, and the following detailed example will illustrate how a mismatch between the components impacts organisational learning and the capacity for such learning.

In the first of the three case studies, a primary goal was to reduce the time taken to change over an automated brazing carousel from one product to another as part of the company's progress to Lean Production. In the current state of affairs a changeover took from 1-1/2 to 2 h and it was felt that this amount of downtime could not be afforded during regular shifts, so production was batched and changeovers were normally conducted on third shift when no production was scheduled. With the goal of setup time reduction in mind, the training and

orientation portion of the Kaizen event featured selected training materials on the subject of quick changes so that the participants were exposed to a wide variety of techniques that they might be able to apply.

One of the lead hands who routinely performed the changeovers was a member of the team, which was significant because he would have to implement the subsequent revisions. What happened was that as long as the efforts were directed towards simply reducing the amount of work required for changeovers, there was little debate and the proposed improvements were readily accepted and implemented. A great deal of creativity was demonstrated, and some of the changes were quite ingenious.

However, it is also possible to reduce the actual downtime by other means such as pre-staging some of the work before the shutdown or deferring some of the clean-up until after it. It is also possible to add workers while the machine is stopped even though they may not be performing as efficiently as they could if they weren't getting in each other's way (Hall 1983). When such suggestions created the possibility of further decreasing machine downtime at the expense of adding man-hours, the reaction was interesting. These proposals were not challenged; they were simply ignored. In this event it was proposed and eventually (somewhat reluctantly) agreed to build some pre-staging racks for setting up the new fixtures in advance, but very little progress was made on implementing this change. It was observed that, somehow, those working on it seemed to keep finding other tasks to have higher priority. When asked about this, and the idea of bringing in a second person to speed the changeover, the individuals most directly involved didn't really want to challenge the training materials, but made it clear that these ideas didn't make sense to them in their situation.

Upon further questioning, the participants demonstrated that they were motivated primarily by an interest in minimising the hours booked for fixture changes. This of course is not the point of setup reduction under Lean Production—these workers were solving a different problem. Although apparent progress was made on the performative elements of the problem, it was the hidden ostensive component that ultimately limited the organisation's ability to “learn” in the sense of improving situated performance. This is the point made by Barnard (1938) who wrote that: “An intelligent person will deny the authority of that ... which contradicts the purpose of the effort as *he* understands it” (p. 166, emphasis in the original).

Although the reported results from this event were good (50 % reduction in setup times), it represented a failure of sorts in the organisational learning process. What is interesting about this failure is that it occurred in a setting that should have been highly conducive to learning: training was provided, resources were made available as needed and there was visible evidence of upper management support; all things commonly cited as success factors. The team members had all the tools they needed and they had an adequate understanding of the principles of Lean, yet they had not adopted a truly Lean perspective; instead they were achieving what Emiliani (2007) refers to as “False Lean”.

In examining the reasons for this, it became apparent that knowledge of performative aspects of business processes is cumulative. That is to say, an organisation can be more or less knowledgeable on a continuum and it is reasonable to accept that the training referred to above had made the organisation more knowledgeable with respect to the performative aspects of Lean. The same, however, is not true of the ostensive elements. While a person or organisation may simultaneously know how to perform multiple tasks, they cannot simultaneously hold different opinions about which is the best under the circumstances. In that sense, the ostensive component of organisational knowledge is something of an either/or proposition. Although training had been provided in the ostensive aspects of Lean Production, and the participants understood it well enough, it had failed to displace the existing mind-set of those charged with implementation. We can state this as our second broad proposition:

*P2: If organisational learning requires changes to a goal or the path to achieving that goal, then displacement of the previous goal or approach is an essential element of the learning.*

It is clear that when we talk about a goal or an approach to a goal, we are discussing the ostensive component of organisational knowledge. Goals and approaches are somewhat interchangeable depending on the level we are looking at within the organisation. For example, a CEO may decide that to achieve a goal of increased shareholder value, it is better to switch from a strategy of rigorous cost-cutting to one of radical product innovation. This represents simply a change in approach and while there may be significant operational challenges, it is not all that difficult conceptually. At the shop floor however, this is a fundamental change of goal and requires a significant reconceptualisation of the purpose of the enterprise. It is not surprising that change initiatives often stall when they reach this level.

The concept of displacement corresponds well to the term *unlearning*, the importance of which has been noted by Weick (1979), Schein (1993) and Pentland (1995). Recognising the need for displacement is much easier than prescribing how to achieve it. What does seem apparent is that failure of the present logic is not sufficient since individuals will remain committed to their plans of action in the face of substantial evidence of their lack of utility (Mitroff and Mason 1974). Some simple examples from our cases illustrated this. There were some problems that had to be dealt with along the way (two examples: parts sticking to fixtures, brazing strips breaking at the feeders) where the proposed solutions depended on beliefs about the root causes. In general, failure of the proposed solutions did not cause the operators to re-assess their beliefs about root causes. That only occurred when they were faced with clear evidence that an approach based on a different causal mechanism did work.

Unfortunately, for more important issues where the payoff is in the future and not necessarily deterministic (for example, success of a market strategy may depend on a competitor's response), clear evidence is hard to come by. In these situations we see that the ostensive component of organisational knowledge corresponds well to Plato's description of knowledge as "justified belief" (Nonaka and Takeuchi 1995).

When organisational learning requires change to the ostensive component (the belief), it is not only necessary that the new approach be justified, it must be better justified than what it is to replace. This forces us to confront the basis on which existing beliefs are justified and is the reason that this article positions organisational learning as a fundamentally behavioural issue.

These cases provided no very good examples of how this displacement process might actually work, in spite of the fact that efforts were made during training to instil a new “justified belief”. All of these events were focused on making progress towards Lean Production, and to that end, each event included training about the philosophy and benefits of Lean. Most participants seemed to understand this material and could articulate it, but subsequent conversations made it very clear that this training had done nothing to change their prevailing mind-sets. To capture this in an abbreviated way: in all three companies, it seemed that floor-level workers had an intuitive understanding that their mission was to maximise output per unit of their time. Nothing in the training had caused them to reject this view in favour of other approaches, as for example when a Lean production schedule with smaller lot sizes would require them to work harder and actually produce less. This clearly did not make sense in their world. To some extent that serves as evidence for the difficulty involved when there is a lack of clear empirical support for displacement of the prevailing approach; after all, what they were doing was apparently adequate yesterday, so why was it now unacceptable?

## 5 Implications for Practice

Managers have a clear interest in improving situated performance and are therefore interested in improvements to the state of organisational knowledge through a learning process. What this study has shown is that there are two distinct components of this organisational knowledge that must be addressed simultaneously if the expected result is to be achieved. This is made more difficult by the fact that the two components cannot be addressed in the same way. Improvements in performative capability are readily achieved by conventional means: training, resources, management support, metrics and incentives. These methods however, are ineffective in achieving the displacement necessary for improvements in the ostensive component of organisational knowledge. A full examination of the means by which these “justified beliefs” persist or are changed is beyond the scope of this study, but the evidence is sufficient to suggest that there is a strong cultural component to beliefs and that they change slowly—if at all.

This tends to be particularly true if the existing beliefs have been validated by years of apparent success as was the case in all of the companies in this study. Under these conditions it seemingly requires a crisis in order to displace the current beliefs and open the door to the desired change. It has even been suggested that crises can be manufactured for this purpose (Kim 1998). A less draconian approach is to assume that, at least during a transition period, that one simply

cannot count on the rank and file to “get it” and to interpret the problem correctly or to apply their expertise in the right way. This creates the somewhat counter-intuitive recommendation that we should focus less on results achieved and more on the activities and behaviours involved. A good example of this was observed in the third company studied. This company was judged to be the furthest along of the three in terms of Lean philosophy. One of the policies that had been put in place was that they would build exactly to order, no more and no less. On some days that meant setting up to produce one unit of a product variant. In spite of their Lean training, many of the employees were strongly opposed to this policy; they did not see how it could be worth the trouble involved and wanted very much to be able to run larger batches. It may or may not have been possible to change those beliefs over time, but in the short term the problem was avoided by telling people how to do their jobs and not allowing them to proceed in what they might have thought was the best manner.

This need to focus on behaviour rather than results is rooted in the fact that metrics tend to mask the underlying assumptions and can be manipulated to show results that are at odds with reality (Melnyk et al. 2010). This is essentially what happened in our setup time reduction example: what appeared to be excellent results masked the fact that the problem was not really being approached in the right way. As in that case, it is often possible to achieve a given result in a variety of ways or with different trade-offs. Naturally it is hoped and expected that appropriate choices will be made, but these are rarely spelled out. Instead, the use of metrics and incentives in performance management is heavily dependent on what (Hanson et al. 2011) called *informal alignment*—the extent to which employees have absorbed the prevailing perception of the problem and the appropriate range of options for addressing it. Attempting to change direction by changing the metrics is apt to result in behaviour that is at cross-purposes with the intended change.

To summarise, a manager attempting to foster organisational learning must develop an awareness of the two distinct components of organisational knowledge that must be addressed. This would be followed by an assessment of the gap that exists between the current and the desired states—on both dimensions. Where a gap exists in the ostensive aspect, a choice must be made whether to try to narrow the gap by changing minds—being mindful of the displacement issue and the cultural basis of justified belief—or to bridge the gap by dictating behaviour as was done in our example. Our evidence suggests that the latter is likely to be more effective, and this finding is echoed by Etlie and Rosenthal (2008).

## 6 Implications for Research

The problems facing the researcher are in many ways similar to those faced by management in the sense that both are, to a degree, outsiders when it comes to understanding what is really going on in an organisation and must rely on



standardised measures to tell the story. The procedural or performative aspects of organisational knowledge are much easier to observe and catalogue than the ostensive aspects and as a result researchers are often in the position of trying to explain situated performance as a dependent variable on the basis of the performative elements as independent variables. In light of the above discussion and the cited examples from the cases, it is clear that this creates a major missing variable problem. This has been very apparent in the work on Lean Production, where the procedural elements of Lean have proved to be of limited value in predicting the ultimate success of the initiative (Oliver et al. 1996; Lewis 2000).

In essence, this is a problem of interpreting variance in observed behaviour (Pentland 2003). Armed with the two-part view of organisational knowledge, we can see that variance can be good or bad: it can be the result of intelligent fine-tuning of the performance to suit the circumstances, or it can be an inability to follow best practices. As a result, we cannot know whether certain observable characteristics should be predictive of success or failure until we know something about the motivation behind them—the ostensive component of organisational knowledge. Absent that knowledge, a data set that contains examples of both types of variance can be expected to produce equivocal or non-significant results.

Unfortunately, this component is not readily accessible to conventional research techniques, particularly since the participants themselves may not be able to fully articulate the logic behind their actions (Cohen and Bacdayan 1994). Furthermore, there is no reason to assume that the logic that exists at the researcher's point of contact (typically the person filling out a survey or being interviewed) is the same as that where the work is being done (Rothaermel and Hess 2007). This state of affairs does not lead to easy solutions, but does suggest some recommendations. The first, obviously, is awareness. When we recognise that situated performance is dependent on two very distinct components of organisational knowledge, we are better equipped to explain what we see. Secondly, we must design our research methods to ask the right questions. While we may not get total clarity, we must start to see organisations as collections of individuals acting within a cultural context whose characteristics are at least notionally discoverable. Finally, these studies have made it clear that if we want to better understand how organisations learn and function, we are going to have to get more deeply embedded in them. Ethnography is our guide here and teaches us that cultures are not readily understood from the outside.

## 7 Conclusion

The value of knowledge as an asset of an organisation is such that a great deal of attention has been paid to its “management”. Unfortunately, the language of knowledge management has created an image of organisational knowledge as having an independent existence such that it can be stored, retrieved, transmitted, absorbed and replicated. A pragmatic view of knowledge (situated performance)

suggests that none of these terms are strictly applicable and that knowledge consists of two parts that are incommensurate, but must both exist in the same place at the same time.

As a result, we cannot manage knowledge directly; rather we must manage its component parts in order to create new and unique instances of organisational knowledge as required. Managing the performative aspects is quite well understood, and many examples of knowledge management in practice do just that through training and documentation. Managing the ostensive component is not only more difficult, as these cases have shown, but ultimately more important. Because the ostensive component provides the core of “justified belief”, the mechanics of justification must be understood. Although this study has barely scratched the surface in this respect, it was clear that justified belief is rooted in, and specific to the cultural setting—to the extent that an individual might legitimately hold different beliefs in different settings. This promises to be the next frontier in knowledge management—one that is not based on artificial intelligence and information technology, but instead in the behavioural sciences.

## References

- A. Amin, P. Cohendet, Organisational learning and governance through embedded practices. *J. Manag. Gov.* **4**(1–2), 93 (2000)
- S.D. Amundson, Relationships between theory-driven empirical research in operations management and other disciplines. *J. Oper. Manag.* **16**(4), 341–359 (1998)
- C.A. Anderson, J.J. Lindsay et al., Research in the psychological laboratory: truth or triviality? *Curr. Dir. Psychol. Sci.* **8**(1), 3–9 (1999)
- C. Argyris, D.A. Schön, *Organizational Learning: A Theory of Action Perspective* (Addison-Wesley Publishing Company, Reading, 1978)
- C.I. Barnard, *The Functions of the Executive* (Harvard University Press, Cambridge, 1938)
- J. Barney, Firm resources and sustained competitive advantage. *J. Manag.* **17**(1), 99 (1991)
- C.A. Bartel, R. Garud, in *Narrative Knowledge in Action: Adaptive Abduction as a Mechanism for Knowledge Creation and Exchange in Organizations*, ed by M. Easterby-Smith, M.A. Lyles. The Blackwell Handbook of Organizational Learning and Knowledge Management (Blackwell Publishing, Oxford, 2003)
- K.A. Bates, E.J. Flynn, Innovation history and competitive advantage: a resource-based view analysis of manufacturing technology innovations. *Acad. Manage. Proc.* **1995**(1), 235–239 (1995). doi:[10.5465/ambpp.1995.17536502](https://doi.org/10.5465/ambpp.1995.17536502)
- R.A. Bettis, M.A. Hitt, The new competitive landscape. *Strateg. Manag. J.* **16**, 7 (1995)
- M.D. Cohen, Individual learning and organizational routine: emerging connections. *Organ. Sci.* **2**(1), 135–139 (1991)
- M.D. Cohen, P. Bacdayan, Organizational routines are stored as procedural memory: evidence from a laboratory study. *Organ. Sci.* **5**(4), 554–568 (1994)
- W.M. Cohen, D.A. Levinthal, Absorptive capacity: a new perspective on learning and innovation. *Adm. Sci. Q.* **35**, 128–152 (1990)
- D.J. Collis, C.A. Montgomery, *Corporate Strategy: A Resource-Based Approach* (Irwin/McGraw-Hill, Boston, 1998)
- R.M. Cyert, P. Kumar et al., Information, market imperfections and strategy. *Strateg. Manag. J.* (1986–1998) **14**(SPECIAL ISSUE), 47 (1993)

- R.M. Cyert, J.G. March, *A Behavioral Theory of the Firm* (Prentice-Hall, Englewood Cliffs, 1963)
- T.H. Davenport, L. Prusak, *Working Knowledge: How Organizations Manage What They Know* (Harvard Business School Press, Boston, 1998)
- P. Dennis, *Lean Production Simplified: A Plain Language Guide to the World's Most Powerful Production System* (Productivity Press, New York, 2002)
- J. Dewey, *Democracy and Education: An Introduction to the Philosophy of Education* (Collier-MacMillan, London, 1916)
- J. Dewey, A.F. Bentley, *Knowing and the Known* (Beacon Press, Boston, 1949)
- P.F. Drucker, *Post Capitalist Society* (Harper Business, New York, 1993)
- A.C. Edmondson, A.B. Winslow et al., Learning how and learning what: effects of tacit and codified knowledge on performance improvement following technology adoption. *Decis. Sci.* **34**(2), 197 (2003)
- L.M. Ellram, The use of the case study method in logistics research. *J. Bus. Logist.* **17**(2), 93–138 (1996)
- M.L. Emiliani, *Real Lean: Understanding the Lean Management System* (The Center for Lean Business Management, LLC, Wethersfield, 2007)
- J.E. Ettl, S.R. Rosenthal, Service innovation in manufacturing (Final Report: National Science Foundation Grant No.0453694) (2008)
- M.S. Feldman, B.T. Pentland, Reconceptualizing organizational routines as a source of flexibility and change. *Adm. Sci. Q.* **48**(1), 94 (2003)
- M.S. Feldman, B.T. Pentland, Issues in empirical field studies of organizational routines. Working Paper (2005)
- B.B. Flynn, S. Sakakibara et al., Relationship between JIT and TQM: practices and performance. *Acad. Manag. J.* **38**(5), 1325–1360 (1995)
- R.M. Grant, Toward a knowledge-based theory of the firm. *Strateg. Manag. J.* (1986–1998) **17**(Winter Special Issue), 109 (1996)
- R.W. Hall, *Zero Inventories* (Dow Jones-Irwin, Homewood, 1983)
- J.D. Hanson, S.A. Melnyk et al., Defining and measuring alignment in performance management. *Int. J. Oper. Prod. Manag.* **31**(10), 1115–1139 (2011)
- B. Hedberg, *How Organizations Learn and Unlearn*, ed. by P.C. Nystrom, W.H. Starbuck. Handbook of Organizational Design (Oxford University Press, Oxford, 1981), pp. 3–27
- R. Henderson, I. Cockburn, Measuring competence? Exploring firm effects in pharmaceutical research. *Strateg. Manag. J.* (1986–1998) **15**(SPECIAL ISSUE), 63 (1994)
- B.A. Hoey, What is ethnography? (2011), [http://www.brianhoey.com/General%20Site/general\\_defn-ethnography.htm](http://www.brianhoey.com/General%20Site/general_defn-ethnography.htm). Accessed 15 June 2011
- S.D. Hunt, *Modern Marketing Theory: Critical Issues in the Philosophy of Marketing Science* (South-Western Publishing Co, Cincinnati, 1991)
- A. Kaplan, *The Conduct of Inquiry, Methodology for Behavioral Science* (Chandler Publishing, San Francisco, 1964)
- L. Kim, Crisis construction and organizational learning: capability building in catching-up at Hyundai motor. *Organ. Sci.* **9**(4), 506–521 (1998)
- B. Kogut, U. Zander, What firms do? Coordination, identity, and learning. *Organ. Sci.* **7**(5), 502 (1996)
- D. Leonard-Barton, Core capabilities and core rigidities: a paradox in managing new product development. *Strateg. Manag. J.* **13**(SPECIAL ISSUE), 111 (1992)
- B. Levitt, J.G. March, Organizational learning. *Ann. Rev. Sociol.* **14**, 319–340 (1988)
- M.A. Lewis, Lean production and sustainable competitive advantage. *Int. J. Oper. Prod. Manag.* **20**(8), 959–978 (2000)
- J.G. March, H.A. Simon, *Organizations* (Wiley, New York, 1958)
- J. Martin, *Organizational Culture: Mapping the Terrain* (Sage, Thousand Oaks, 2002)
- S.A. Melnyk, R.J. Calantone et al., Short-term action in pursuit of long-term improvements: introducing Kaizen events. *Prod. Inventory Manag. J.* **39**(4), 69–76 (1998)

- S.A. Melnyk, J.D. Hanson et al., Hitting the target but missing the point: resolving the paradox of strategic transition. *Long Range Plan.* **43**(4), (2010)
- M.B. Miles, A.M. Huberman, *Qualitative Data Analysis: An Expanded Sourcebook* (Sage, Thousand Oaks, 1994)
- D. Miller, Configurations revisited. *Strateg. Manag. J.* (1986–1998) **17**(7), 505 (1996)
- I.I. Mitroff, R.O. Mason, On evaluating the scientific contribution of the Apollo moon missions via information theory: a study of the scientist–scientist relationship. *Manag. Sci.* **20**(12), 1501–1513 (1974)
- A.S. Mukherjee, M.A. Lapre et al., Knowledge driven quality improvement. *Manag. Sci.* **44**(11), S35 (1998)
- J. Nahapiet, S. Ghoshal, Social capital, intellectual capital, and the organizational advantage. *Acad. Manag. Rev.* **23**(2), 242 (1998)
- R.R. Nelson, Why do firms differ, and how does it matter? *Strateg. Manag. J.* (1986–1998) **12**(SPECIAL ISSUE), 61 (1991)
- I. Nonaka, A dynamic theory of organizational knowledge creation. *Organ. Sci.* **5**(1), 14 (1994)
- I. Nonaka, H. Takeuchi, *The Knowledge-Creating Company* (Oxford University Press, New York, 1995)
- T. Ohno, *Toyota Production System: Beyond Large-Scale Production* (Productivity Press, New York, 1988)
- N. Oliver, R. Delbridge et al., Lean production practices: international comparisons in the auto components industry. *Br. J. Manag.* **7**(1), 29–44 (1996)
- W.J. Orlikowski, Knowing in practice: enacting a collective capability in distributed organizing. *Organ. Sci.* **13**(3), 249–273 (2002)
- E.T. Penrose, *The Theory of the Growth of the Firm* (Oxford University Press, Oxford, 1959)
- B.T. Pentland, Organizing moves in software support hot lines. *Adm. Sci. Q.* **37**(1992), 527–548 (1992)
- B.T. Pentland, Grammatical models of organizational processes. *Organ. Sci.* **6**(5), 541 (1995)
- B.T. Pentland, Sequential variety in work processes. *Organ. Sci.* **14**(5), 528 (2003)
- B.T. Pentland, M.S. Feldman, Organizational routines as a unit of analysis. *Ind. Corp. Change* **14**(5), 793–815 (2005)
- M. Polanyi, *Personal Knowledge: Towards a Post-Critical Philosophy* (University of Chicago Press, Chicago, 1958)
- C.K. Prahalad, G. Hamel, The core competence of the corporation. *Harv Bus Rev* **68**(3), 79 (1990)
- E. Rosch, B.B. Lloyd (eds.), *Cognition and Categorization* (Lawrence Erlbaum Associates, Hillsdale, 1978)
- F.T. Rothaermel, A.M. Hess, Building dynamic capabilities: innovation driven by individual, firm, and network-level effects. *Organ. Sci.* **18**(6), 898–921 (2007)
- R. Sabherwal, I. Becerra-Fernandez, An empirical study of the effect of knowledge management processes at individual, group, and organizational levels\*. *Decis. Sci.* **34**(2), 225 (2003)
- E.H. Schein, How can organizations learn faster? The challenge of entering the green room. *Sloan Manag. Rev.* **34**, 85–92 (1993)
- J.-C. Spender, Making knowledge the basis of a dynamic theory of the firm. *Strateg. Manag. J.* (1986–1998) **17**(Winter Special Issue), 45 (1996)
- W.H. Starbuck, Learning by knowledge-intensive firms. *J. Manag. Stud.* **29**(6), 713 (1992)
- D.J. Teece, G. Pisano et al., Dynamic capabilities and strategic management. *Strateg. Manag. J.* (1986–1998) **18**(7), 509 (1997)
- J.P. Walsh, Managerial and organizational cognition: notes from a trip down memory lane. *Organ. Sci.* **6**(3), 280 (1995)
- K.E. Weick, *The Social Psychology of Organizing* (McGraw-Hill, New York, 1979)
- B. Wernerfelt, A resource-based view of the firm. *Strateg. Manag. J.* **5**(2), 171 (1984)
- O.E. Williamson, The economics of organization: the transaction cost approach. *Am. J. Sociol.* **87**(3), 548–577 (1981)
- R.K. Yin, *Case Study Research Design and Methods* (Sage Publications, Thousand Oaks, 1994)