

Transferring, Translating, and Transforming: An Integrative Framework for Managing Knowledge Across Boundaries

Paul R. Carlile

Information Systems Department, School of Management, Boston University, 595 Commonwealth Avenue,
Boston, Massachusetts 02215, carlile@bu.edu

The paper examines managing knowledge across boundaries in settings where innovation is desired. Innovation is a useful context because it allows us to explore the negative consequences of the path-dependent nature of knowledge. A framework is developed that describes three progressively complex boundaries—syntactic, semantic, and pragmatic—and three progressively complex processes—transfer, translation, and transformation. The framework is used to specify the practical and political mismatches that occur when innovation is desired and how this relates to the common knowledge that actors use to share and assess each other's domain-specific knowledge. The development and use of a collaborative engineering tool in the early stages of a vehicle's development is presented to illustrate the conceptual and prescriptive value of the framework. The implication of this framework on key topics in the organization theory and strategy literatures is then discussed.

Key words: knowledge; innovation; path dependence; power; boundary management; product development

1. Introduction

Dorothy Leonard's (1995) statement that most innovation happens at the boundaries between disciplines or specializations tells us that working across boundaries is a key ingredient of competitive advantage, but also why innovation proves so difficult to create and maintain. The growing research on knowledge in organizations underscores this challenge by recognizing first the "knowledge boundaries" (Brown and Duguid 2001) between specialized domains and second that knowledge is "both a source of and a barrier to innovation" (Carlile 2002, p. 442). In this paper the topic of boundaries in organization theory will be reexamined from a perspective of managing knowledge across boundaries in settings where innovation is desired.

This focus on innovation will help to extend the conversation of boundary management from its classic starting point of information processing (Lawrence and Lorsch 1967, Galbraith 1973) and beyond its more contemporary focus on coordination (Malone and Crowston 1994, Gittell 2001). This examination will provide analytic descriptions of the varying circumstances possible at boundaries and the processes involved in managing knowledge across them. This effort will help resolve the incompatibility between three different perspectives of boundaries: an information processing approach that focuses on knowledge as a thing to store and retrieve, an interpretive approach that emphasizes the importance of a common meaning to share knowledge between actors, and a political approach that acknowledges how different interests impede knowledge sharing. The analytic effort to integrate these often incompatible perspectives sets up the practical purpose of the paper—to describe the

different processes required at each type of boundary to effectively manage knowledge.

To do this, a framework is developed that describes three progressively complex boundaries—syntactic, semantic, and pragmatic (see Shannon and Weaver 1949) and three progressively complex processes—transfer, translation, and transformation. A focus on the effectiveness of managing knowledge across boundaries clarifies that the relationship between actors is one where they not only share their knowledge, but also assess each other's knowledge. Implicit in the effectiveness of this effort is the existence of a common knowledge¹ that actors use to share and assess each other's domain-specific knowledge. Acknowledging both domain-specific knowledge and common knowledge at a boundary provides a useful distinction to better understand the challenges as actors try to work across domains when innovation is desired.

The empirical context used to illustrate this framework and understand these challenges will be the boundaries that exist among specialized domains in a new product-development setting. These boundaries are especially challenging in the early development stages where the impact of new requirements is hard to determine. The following description offers a glimpse into the challenges at this initial stage and introduces the case that will be used to illustrate the conceptual and prescriptive value of the framework.

Beta Motors, like any product-development firm, faces significant challenges in the early design stage of a new vehicle. The functional groups involved at this early design stage are vehicle styling, engine/power train,

climate control, and safety. The early stage represents a unique opportunity to specify up front the overall design constraints of the vehicle. If done well, it results in fewer downstream problems such as conflicts between engineering groups, launch delays, costly rework on the line, and even warranty issues.

Since the 1960s, the most common engineering tool used to share and assess knowledge across these groups was the use of a variety of “clay models” of the vehicle’s design. However, when increasing quality and time-to-market pressures occurred in the 1980s, conflicts between groups, cost overruns, and launch delays also steadily increased. By the early 1990s, Beta Motors spent significant resources to develop better methods to manage knowledge at this early design stage. In 1995 the deployment of a new simulation tool proved extremely successful—yet on its fourth deployment, expensive delays and quality problems once again arose.

The framework developed will help explain why the use of the clay models and the later deployment of the simulation tool failed as a common knowledge for the groups (actors) involved to adequately share and assess each other’s domain-specific knowledge.

The first section of the paper is primarily a conceptual effort to develop the framework. Three relational properties of knowledge at a boundary are introduced to describe the varying circumstances that can exist when managing it. The article then describes three types of boundaries and their connection to the existing literature on knowledge and boundaries and how, together, they address the increasingly complex circumstances possible at a boundary. In the second section, the case is developed more fully, with the development of the modeling tool, its improvement over the use of clay models, yet its failure in its fourth deployment described. The third section focuses on how the presence of novelty required for innovation clashes with the path-dependent tendency of knowledge, creating relations between actors that limit their ability to effectively share and assess each other’s knowledge. The article then discusses the implications this integrative framework has on key topics in organization theory.

2. Describing the Relative Complexity of a Boundary

To integrate different approaches to managing boundaries in organization theory, it is important to describe the potential range of circumstances or the relative complexity at a given boundary. To develop the framework, the three following properties of knowledge at a boundary will be discussed: difference, dependence, and novelty (Carlile and Reberich 2003). Difference in knowledge refers to a difference in the amount of knowledge accumulated, for example, the novice-expert distinction found in Schank and Abelson (1977) and Hinds (1999) and/or the difference in the type of domain-

specific knowledge accumulated, such as specialization in different problem-solving domains found in Weber (1924/1947) between actors. Creating a complex product or service often requires differences in the amount and type of knowledge. This in turn creates differences in levels of experience, terminologies, tools, and incentives that are unique to each specialized domain. At Beta Motors actors in vehicle styling, engine/power train, climate control, and safety groups specialize in different kinds of engineering work that demand different types of knowledge and responsibilities.

This specialization of knowledge goes deeper than just the actor’s role or identity. The styling group strives to create a distinctive vehicle in comparison to the competition. The engine/power train group tries to develop the most efficient but powerful engine possible. Climate control needs to adequately cool the vehicle in summer and heat it in winter. Safety must ensure that the vehicle is as safe as possible for its occupants, as well as for the occupants of another vehicle with which it may collide. For this reason knowledge is not only localized but also invested within a given practice. Because knowledge takes investment—time and resources to acquire—it should be seen as “at stake,” indicating the significant costs associated with giving it up and acquiring different knowledge (Carlile 2002). As difference in the amount and/or type of domain-specific knowledge increases between actors, the amount of effort required to adequately share and assess each other’s knowledge also increases.

The second relational property of knowledge at a boundary is dependence—Without dependence, difference is of no consequence. Examples of dependence can be seen in the relations between coauthors working on a paper, employees on different stages of a product line, or a design engineer and a manufacturing engineer in a product-development setting. Dependence was defined by Litwak and Hylton (1962) as a condition where two entities must take each other into account if they are to meet their goals. Victor and Blackburn (1987) stipulated how the actions of actors determine their individual pay-offs or success, specifying the consequential link between the activities and goals of actors who are dependent on each other. Extending these insights in the context of coordination theory, Malone and Crowston (1994) define coordination as the management of dependence among activities (tasks) and resources (see also Crowston 1997). Substituting domain-specific knowledge for resources, we begin to see some of the complexity that is not always recognized in a coordination theory inspired by artificial intelligence—for if knowledge is different in kind, and not just in degree, then managing dependencies requires the capacity to develop an adequate common knowledge as resources and tasks change.

While Thompson’s (1967) categorization of three types of system-level interdependence (pooled, sequential, and reciprocal) remains an important perspective on

dependence, the focus of this paper is how this management is done. For example, even though sequential interdependence is seen as less complex than reciprocal, the concern is how such dependencies are managed when new circumstances emerge. The approach developed here is more in keeping with Henderson and Clark's (1990) concern that one of the biggest challenges in a complex technology development setting is making visible the architectural dependencies that are of consequence.

At Beta Motors the dependence between vehicle styling and engine/power train recognizes that a bigger engine raises the level of the hood, while the desired look and feel embodied in the drawings constrains the size of the engine. Thus, as the number of dependencies increase between actors, the complexity and the amount of effort required to share and assess knowledge at a boundary also increases. In settings where innovation is required, the coordination of what is known as well as making visible the consequences of dependence that are not currently known are vitally important.

The third relational property of knowledge at a boundary is how novel the circumstances are. In a new product-development setting, the most obvious source of novelty is new customer needs that generate new requirements of the various actors in their specialized domains. This suggests that the most challenging aspect of the relational nature of knowledge at a boundary is that for each actor there is novelty to share with others and novelty to assess from others. A less-obvious source of novelty comes when an actor is unfamiliar with the common knowledge being used to represent the differences and dependencies between domain-specific knowledge. When novelty arises there is often a lack of common knowledge to adequately share and assess domain-specific knowledge at a boundary. A common knowledge could be the use of the English language by actors to communicate and collaborate or more specifically the use of a prototyping methodology. In any case, the common knowledge functions as a boundary object (Star 1989, Carlile 2002), which the actors use to communicate across domains. When novelty is present both the *capacity* of the common knowledge to represent the differences and dependencies now of consequence and the *ability* of the actors involved to use it become important issues.

The word novelty instead of "uncertainty" is deliberately used. Uncertainty is an external characterization (Tsoukas 2001) that all is not known in a given environment, whereas novelty suggests no external vantage point. Novelty underscores the participatory and relational nature of what an actor needs to share and to assess when all is not known. Further, unlike uncertainty, novelty doesn't allow us to take for granted that what is new is easily recognized as something unknown. Actors are susceptible to misrecognizing what is novel as

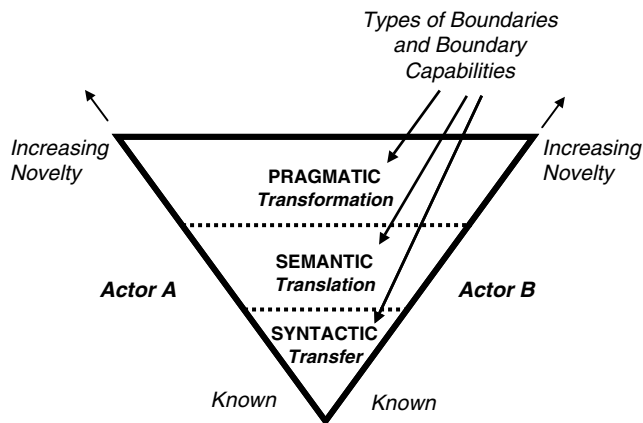
something that is already known (i.e., competency traps as seen in Levitt and March 1988, Martins and Kambil 1999) or discarding what is novel as irrelevant (Perrow 1994, Arrow 1994). Camerer et al. (1989) have referred to this tendency as the "curse of knowledge," which recognizes the difficulty that actors have in abandoning previous knowledge (i.e., knowledge is "at stake"). These issues highlight the challenges that actors face in identifying what is of consequence when novel circumstances arise (Weick et al. 1999). For example, if a previous engine is reused in a vehicle's redesign, there is potentially a significant amount of common knowledge that can be used to share and assess knowledge about the impact of the engine across styling, engine/power train, climate control, and safety. However, when a new engine is used, the amount of common knowledge available across the groups decreases, limiting their ability to represent the differences and dependencies that are now of consequence. So, as novelty increases, the amount of effort required to adequately share and assess knowledge also increases.

We can represent these relational properties of knowledge by imagining a boundary as a vector between at least two actors. It starts at the origin where the differences and dependencies are known; as novelty increases the vector spreads, scaling the increasing complexity and the amount of effort required to manage the boundary. When the relevant differences and dependencies are known, the reuse of the common knowledge has positive effects and the path-dependent nature of knowledge (March 1972) proves beneficial. However, when novelty increases, the path-dependent nature of knowledge has negative effects (Hargadon and Sutton 1997) because the common knowledge used in the past may not have the capacity to represent the novelties now present (Carlile and Reberich 2003). A problematic scenario that often arises is when a powerful actor reuses a common knowledge (i.e., path dependency) that constrains the capacity and ability of other actors to represent the novelty they are facing. Such a mismatch at a boundary puts the first actor in a more powerful position to represent his or her domain-specific knowledge. However, without the development of an adequate common knowledge (i.e., capacity) to represent the novelty now present and/or the ability of the actors to use it, the outcome produced will generate latent and costly problems in the long term.

3. Developing the Integrative Framework

In this section three different approaches to managing boundaries in the organization theory and product-development literatures are discussed and integrated in single framework. Using the image of the vector described above, we scale the relative complexity of the circumstances at a boundary using Shannon and Weaver's (1949) three levels of communication complexity: syntactic, semantic, and pragmatic (see Figure 1). While

Figure 1 An Integrated/3-T Framework for Managing Knowledge Across Boundaries



these three terms also have significant history in the field of linguistics (Cruse 2000), the focus here—on the effectiveness of sharing and assessing knowledge across boundaries—broadens the concern from just the structure, meaning, or use of language, to Shannon and Weaver’s practical concerns about what is required for effective communication across domains. In comparison to Shannon and Weaver, the difference in the use of the three terms here is that although they acknowledged all three levels, their mathematical information theory focused almost entirely at the syntactic level.

3.1. A Syntactic or Information-Processing Boundary: Transferring Knowledge

The most common label used to describe the movement of knowledge in organizations is “knowledge transfer” (Winter 1987, Szulanski 1996, Argote 1999). The concept of “transfer” has its basis in the information-processing approaches to boundaries in organization theory (Lawrence and Lorsch 1967, Galbraith 1973), the roots of which stem from Shannon and Weaver’s mathematical approach to communication and information (1949). The practical strength of their approach is its mathematical capacity to process a syntax (i.e., 0s and 1s in the case of computer technologies) that adequately define the relations between sender and receiver at a boundary. When common lexicon sufficiently specifies the differences and dependencies of consequence at the boundary, the boundary proves “unproblematic”; the primary concern is one of “processing” or transferring knowledge across it. What is not always acknowledged from such a perspective are the stable conditions that allowed a common lexicon to be created and to adequately function as a common knowledge. This failure often leads to underestimating the effort required when those stable conditions change.

Given its technical roots, it is not surprising that an information-processing approach is the dominant view

used to describe managing boundaries in product development (Brown and Eisenhardt 1995) and organization design more generally (Lawrence and Lorsch 1967, Tushman and Nadler 1978). Further, information-processing assumptions are the basis of most technology-based approaches to knowledge management, where the primary focus has been on the storage and retrieval of knowledge (Davenport and Prusak 1998). Simply transferring knowledge, however, proves problematic when novelty arises because the current lexicon is no longer sufficient to represent the differences and dependencies now present. The limitation of an information-processing approach occurs because the processing of a common lexicon is assumed to be always a sufficient common knowledge (Reddy 1979). So while a common lexicon is always necessary, it is not always a sufficient type of common knowledge to share and assess domain-specific knowledge. In the case of Beta Motors, as time-to-market and quality pressures grew, the clay model (the actor’s common lexicon) no longer had the capacity to sufficiently represent the differences and dependencies that now impacted the design of the vehicle and needed to be resolved at an early stage.

3.2. A Semantic or Interpretive Boundary: Translating Knowledge

The transition from a syntactic to a semantic boundary occurs when novelty makes some differences and dependencies unclear or some meanings ambiguous. When new requirements and/or new actors are present, interpretive differences in what a word, measurement, or outcome means limits the effective management of knowledge between actors. Researchers who adopt an interpretive approach recognize how different domains (i.e., thought worlds) naturally generate interpretive differences and so emphasize processes that help create “shared meanings” (Dougherty 1992) or mechanisms “to reconcile discrepancies in meaning” (Nonaka and Takeuchi 1995, p. 67). Research in this vein emphasizes the role of cross-functional teams (Ancona and Caldwell 1992), colocation, and the use of various shared methodologies (i.e., CAD/CAM, cross-functional problem-solving templates) to do this. Others have focused on the role of particular individuals as brokers and translators who enable the flow of knowledge (Allen 1977, Hargadon and Sutton 1997). Further, the communities of practice literature (Lave and Wenger 1991, Brown and Duguid 1991) show that as individuals participate in similar activities, they develop shared meanings (i.e., Orr’s 1996 study of Xerox repairmen).

Many researchers who focus on knowledge at the project or firm level have paid attention to the distinction between tacit and explicit knowledge (Polyani 1966) to recognize the situated and interpretive challenges of moving knowledge across boundaries (Nonaka and Takeuchi 1995, Spender 1996). von Hippel and

Tyre's work (von Hippel 1994, Tyre and von Hippel 1997) calls attention to the stickiness of situated knowledge when trying to move it across different domains. Nonaka's description of the process of "externalization," making tacit knowledge explicit, has been recognized as one of the most critical processes organizations need (Nonaka 1994). All of this research acknowledges the importance of developing common meaning as a way to address interpretive differences across boundaries.

Under some circumstances, however, it is not just a matter of translating different meanings, but of negotiating interests and making trade-offs between actors (Wenger 1998, Brown and Duguid 2001). Given their focus on meaning, perspectives that take an interpretive approach frequently do not specify processes that deal with different interests and their political consequences. For example, Nonaka's process of externalization does not recognize that in the course of making one's knowledge explicit, different interests are often revealed that create barriers to developing shared meanings. Under these circumstances creating common meaning is not possible; what is required is a process in which actors negotiate and are willing to change the knowledge and interests from their own domain.

At Beta Motors, before the simulation tool, clay modeling techniques lacked a syntactic and semantic capacity to represent and discuss the impact of the engine/power train systems on crash-test outcomes. Because of that, it did not allow these groups to identify trade-offs, test potential outcomes, and transform the vehicle's design to avoid costly changes and delays downstream. Today the new simulation tools provide an opportunity for the engine/power train and safety groups to agree about the importance of "crash-test data." Yet what remains is exploring the consequences those results have on each other's work and renegotiating the location of engine and power train systems to improve the crash-test results (i.e., an "acceptable" amount of intrusion into passenger compartment). When different interests arise, developing an adequate common knowledge is a political process of negotiating and defining common interests.

3.3. A Pragmatic or Political Boundary: Transforming Knowledge

The transition from a semantic to a pragmatic boundary arises when the novelty presents results in different interests among actors that have to be resolved. What this recognizes is that when actors have different interests, the dependencies between them are not indifferent (James 1907)—And these pragmatic differences generate costs to the actors involved. Under these circumstances domain-specific knowledge, as well as the common knowledge used, may need to be transformed to effectively share and assess knowledge at the boundary. A focus on pragmatic boundaries recognizes that knowledge is invested in practice and so is "at stake"

for those actors who have developed it (Carlile 2002). When interests are in conflict, the knowledge developed in one domain generates negative consequences in another. Here the costs for any actor are not just the costs of learning about what is new, but also the costs of transforming "current" knowledge being used (i.e., common and domain-specific knowledge). These costs negatively impact the willingness of an actor to make such changes, which helps explain the path-dependent tendency of actors' knowledge despite the presence of novelty.

Research that acknowledges these pragmatic differences frames knowledge processes as "creative abrasion" (Leonard-Barton 1995) and focuses on the negotiation of practice (Brown and Duguid 2001) and the transformation of knowledge (Carlile 2002). In addition to the importance of teams, this work also recognizes the role that shared artifacts and methods play in providing the capacity to negotiate interests and transform knowledge. Boundary objects (Star 1989, Carlile 2002, Bechky 2003) such as drawings, prototypes (Wheelwright and Clark 1995, Schrage 1999), and "trade-off" methodologies (Carlile and Lucas 2003) have proved effective in providing a concrete means of representing different functional interests and facilitating their negotiation and transformation in product-development settings.

For example, in 1990 the engine/power train group at Beta Motors wanted to place their newest, most powerful engine into the vehicle platform. This engine was a breakthrough for the engine group (representing an outcome of a sustained effort over several years) because it produced significant horsepower while still achieving "good" gas mileage. The problem, however, was that the size and the shape of the engine caused the hood to go up much higher than the styling group wanted. Unlike the 1980s when "bulky" trucks were the norm, in the 1990s the market was demanding and competitors were creating increasingly "aerodynamic-looking" vehicles. For the engine/power train group the novelty encountered at this boundary was a smaller engine compartment. For vehicle styling, it was a powerful engine that was too big. Unfortunately the specific consequences of these dependencies were not well understood in the early stages of the vehicle's development. Not only did this generate costly design changes and delays downstream, but the political fallout of these costs drove a large wedge between these two powerful engineering groups. It was this event the led to the final push to develop a new engineering tool to address the costly problems of sharing knowledge across domains.

At a pragmatic boundary actors must be able to represent current and more novel forms of knowledge, learn about their consequences, and transform their domain-specific knowledge accordingly. This knowledge is a transformed mixture of the knowledge determined to still be of value and the knowledge that has been determined of consequence given the novelty present.

3.4. Integrating Different Approaches: A Theoretical Framework

The purpose of discussing these different approaches to boundaries in the product-development and organization theory literatures (see Table 1 for summary) has been to embrace the accurate insights of each while avoiding the incompatibilities that arise when changing circumstances are not specified. In Figure 1 each type of boundary is categorized, showing the relative complexity of a boundary according to the relational properties of knowledge discussed earlier. The framework also identifies different types of processes—transferring, translating, and transforming—associated with each type of boundary. At the bottom of the inverted triangle a syntactic boundary is faced, and when knowledge is transferred according to a common lexicon, domain-specific knowledge can be efficiently managed across the boundary. As novelty arises, new differences and dependencies exist that need to be identified and their consequences understood. At a semantic boundary, a process of learning about and translating domain-specific knowledge establishes common meanings that become adequate for the actors involved to share and assess their knowledge. If a semantic response does not resolve the problem, then a pragmatic boundary is faced. What is now required is negotiating and transforming both the common knowledge and domain-specific knowledge used in the past. Here common interests are developed

that allow actors to address the consequences, differences, and dependencies of each other's domain-specific knowledge.

Although the line between each type of boundary is clearly demarcated in Figure 1, the transition where one ends and another begins is not often easily identified by the actors involved. Further, the purpose of the hierarchical representation in Figure 1 is to recognize that as we move up in complexity, the process or capacity at a more complex boundary still requires the capacities of those below it. For example, an effective transformation process also requires the existence of a common lexicon and meaning. However, these different levels should not be seen like “layers in a communication protocol” as they are in network computing (i.e., the International Standards Organization (ISO) model). In such a mathematically based view of communication the interfaces between increasingly complex layers can be managed through common standards and algorithmic processing. However, this is what breaks down as we move up to semantic and pragmatic levels—and is what concerned Claude Shannon as his mathematic approach was applied above the syntactic level.

At a practical level, this framework can describe the “mismatches” that can occur between the type of boundary faced and the type or capacity of the process used, for example, if a syntactic boundary is faced but the more costly translation and transformation processes are

Table 1 Comparative Summary of Approaches to Sharing and Assessing Knowledge Across Boundaries

	Syntactic boundary: A transfer or information-processing approach	Semantic boundary: A translation or interpretive approach	Pragmatic boundary: A transformation or political approach
Circumstances	Differences and dependencies between actors are known. A common lexicon is developed that is sufficient to share and assess knowledge at a boundary.	Novelty generates some differences and dependencies that are unclear—different interpretations exist. Common meanings are developed to create shared meanings and provide an adequate means of sharing and assessing knowledge at a boundary.	Novelty generates different interests between actors that impede their ability to share and assess knowledge. Common interests are developed to transform knowledge and interests and provide an adequate means of sharing and assessing knowledge at a boundary.
Solutions	<i>Theory:</i> Information processing (Shannon and Weaver 1949, Lawrence and Lorsch 1967)— transferring knowledge <i>Techniques:</i> Syntactic capacity, taxonomies, storage and retrieval technologies.	<i>Theory:</i> Learning (i.e., communities of practice)—creating shared meanings (Dougherty 1992, Nonaka 1994), translating knowledge <i>Techniques:</i> Semantic capacity, cross-functional interactions/teams, boundary spanners/translators	<i>Theory:</i> “Creative abrasion” (Leonard-Barton 1992)—negotiating practice (Brown and Duguid 2001); transforming knowledge (Carlile 2002, Bechky 2003) <i>Techniques:</i> Pragmatic capacity, prototyping and other kinds of boundary objects that can be jointly transformed
Challenges	Increasing capacity to process “more” information (Galbraith 1973) A common lexicon is necessary but not always sufficient to share and assess knowledge across a boundary.	Making tacit knowledge explicit (Polanyi 1966, Nonaka 1994) To create common meanings to share and assess knowledge often requires creating new agreements.	Changing knowledge that is “at stake” (Bourdieu and Wacquant 1992, Carlile 2002) To create common interests to share and assess knowledge requires significant practical and political effort.

developed. A second scenario is when a pragmatic or semantic boundary is faced but only a transfer process is used. This latter scenario is the most strategically dangerous, because the novelty that is not recognized and resolved proves very consequential over time (i.e., see disruptive technologies; Christensen 1997). Implied in this second scenario is the situation in which powerful actors have the tendency to reuse their common knowledge. For those actors the path dependency and perceived efficiency of using a common knowledge means they will not likely recognize the semantic and pragmatic issues that exist at the boundary (i.e., the curse of knowledge). Such path dependency is helpful only when the circumstances at the boundary remain stable.

4. Empirical Case and Approach

This case was developed during a set of follow-up visits to Beta Motors, where there was an ongoing research program studying the utilization of knowledge in product development. Beta Motors is one of the three largest automobile manufacturers in the world. The case was collected at one of its design centers in North America. The focus of the case is the development of a computational fluid dynamic (CFD) tool with a three-dimensional modeling technique and its combined use as a collaborative engineering tool in early design efforts. The case identified is of particular value because it provided both an historical as well as a present-day window into the development of a process of sharing and assessing knowledge at a pragmatic boundary. Further, Beta Motors committed significant resources to carefully deploying the tool across each setting, so it provided a weak but naturalized experiment to understand the particular challenges of establishing and maintaining an adequate process at a pragmatic boundary.

The data for the case were collected over a two-month period, including interviews with seven participants and five follow-up telephone interviews to clarify particular events with key informants. Participants were selected from each of the four engineering groups involved and included the primary developer of the tool and members of his staff as well. The purpose of this case is not to test the integrative framework, but to illustrate its value in connecting the variety of literatures on knowledge and boundaries and describing how the path dependency of knowledge and power create mismatches in effectively managing knowledge across boundaries.

To present the material from this case, I first describe how the tool was developed and then “successfully” used on the B-150 vehicle. Next, I discuss the characteristics that define an effective boundary capability at a pragmatic boundary by linking these characteristics to the framework developed earlier. Finally, I discuss the “failed” use of the tool on the B-100 vehicle and summarize what can be learned from this in terms of how mismatches arise at a boundary.

4.1. Developing and Using the CFD Tool

Given the launch delays and downstream cost overruns that Beta Motors had been experiencing since the early 1980s, the head of the engineering group was looking for a better way to manage the early stages of a vehicle’s design. As a member of the advanced vehicle design group, Bill Knox was asked to develop a tool to make communication and problem solving more effective across the four major groups—vehicle styling, engine/power train, climate control, and safety. What made this task particularly difficult was that the knowledge developed by these different groups had to be invested in meeting the requirements within their specialized domains—yet it was dependent on others to generate an effective joint outcome. For example, styling wants to create an aesthetically distinctive vehicle design that differentiates itself from competitors’ current products, which in turn defines the shape of the overall look and feel of the vehicle. The engine/power train group has horsepower requirements and fuel economy constraints that affect the weight of the vehicle and slope of the hood. The climate control group has to make sure that, given the engine size, the vehicle can stay cool in the summer and warm in the winter with a grill size that doesn’t compromise aesthetics. The safety group has concerns about the placement of bumpers and the location of the engine to limit collision damage to the car and other vehicles involved. For Bill, “knowing the limitations of a design is about understanding the interdependencies across the requirements of the groups involved.”

With the increasing time-to-market and quality pressures in the 1980s, a clay model could not adequately represent differences and dependencies that now needed to be identified and resolved at this early stage. However, because clay models were the primary tools used by the powerful vehicle styling group and were developed early in the redesign of a vehicle, they remained the dominant tool used to shape the design of the car. This was not only due to the engineering culture at Beta Motors, but also reflects that clay models represented the look and feel important to distinguish the design in relation to the look and feel of the competitor’s vehicle. For Bill Knox, what was missing was a tool that the groups could use to represent “the major constraints that would go into defining the overall ‘design space’ of a vehicle.” As a specialist in aerodynamics, Bill had used CFD techniques to assess the aerodynamics and fuel economy of many vehicles. His thought was to use these techniques to represent these cross-disciplinary requirements and “improve joint problem solving.” To do this, he also added a three-dimensional modeling technique to represent the “overall skin” to assess the impact of different components on the vehicle’s shape. At a technical level, he believed that representing the “skin” of the vehicle and measuring its “drag coefficient” trade-offs across the different functional domains could be better managed.

Bill often made reference to a “firemen’s tarp” as the collaborative image that he saw the engineering tool playing. A firemen’s tarp was used in the early twentieth century to catch individuals jumping from burning buildings. Bill always stated that the effectiveness of the tarp came from three things. First, it had to be made of a strong material and fashioned in a way that allowed each fireman to easily hold and use it. Second, it needed to be held by several firemen pulling as hard as they could in different directions to safely break an individual’s fall. And third, the firemen had to constantly look up at the individual and make adjustments to ensure the individual jumping landed safely in the middle of the tarp. Bill used this image as a guide as he set out to create an engineering tool that would establish a “shared way for each group to pull hard in their own direction and still make good trade-offs.”

Over a period of four months Bill and four other engineers consulted each group to understand the design requirements faced within each specialized domain. This information provided the technical basis for how each group’s requirements would be specified using the modeling tool. They also created a format where data such as size, shape, movement, material, heat, and other tolerances could be specified. Bill and his groups also spent time meeting with each group to check if the model did indeed reflect their technical requirements, as well as whether it had the capacity to test their intentions as they learned more. What was most difficult to identify and measure were the dependencies across each engineering domain. For example, what are the consequences of putting in a larger engine on the shape, weight, cooling requirements, gas mileage, and crash-test outcomes? At this early stage they used the “skin” of the car and its “drag coefficient” as measures of dependence to represent and assess the knowledge being used to generate the vehicle’s design. Although Bill and his group of engineers were careful to use the language of each engineering group, they also had to develop a language that was familiar to all the groups involved.

Once Bill and his group felt comfortable that they were able to represent the design requirements of each group, they spent the next two months teaching individuals from each of the four groups how to use the tool, developing a common language, and making changes to the tool that reflected concerns. The biggest challenge they faced was not technical; it was getting all of the groups involved to recognize the benefits of the tool despite some of the control that each group felt it was giving up (i.e., discretion, Barley 1986)—with the historically more powerful groups feeling like they were giving up the most control (i.e., vehicle styling, engine/power train).

After seven months of developing and fine tuning, the CFD modeling tool was given to the B-150 redesign team. The effectiveness of the tool was significant, both

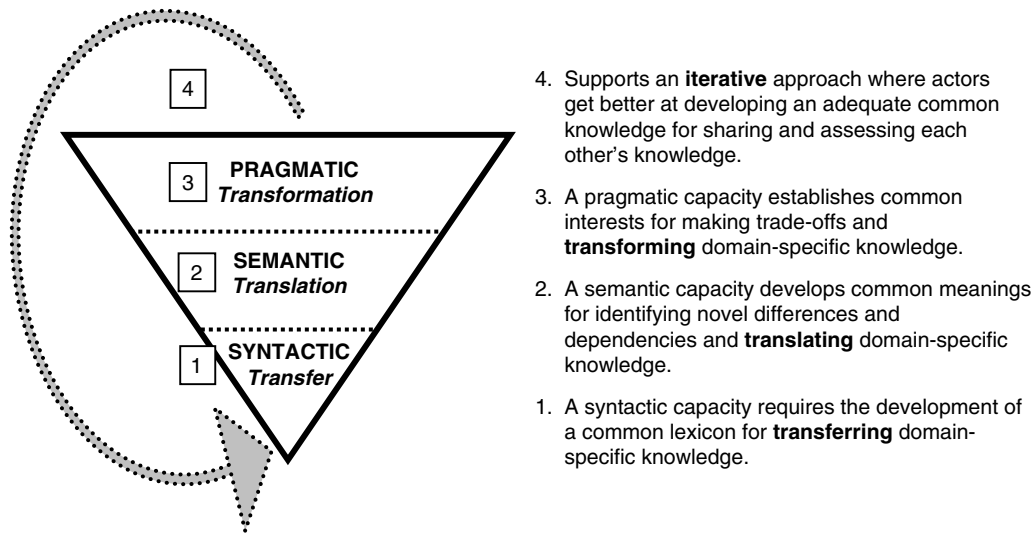
in terms of reducing engineering time by 30% and costs by 50% at the early stage, but more importantly the downstream implementation of the B-150’s redesign avoided the major rework costs and launch delays that had occurred on its last two redesign efforts. Over the next year, the CFD modeling tool was used in three other redesign efforts. Two had similar successes to the B-150; the last, the B-100, had similar savings in up-front engineering time and costs, but when the design was implemented, significant problems were identified downstream. The outcome was that launch delays and costs were as high as previous redesign efforts on the B-100.

4.2. The Characteristics of a Process at a Pragmatic Boundary

To explain the different outcomes of the CFD in the B-150 and the B-100, four characteristics describe the capability required at a pragmatic boundary (see Figure 2). The first characteristic is the development of a common lexicon that actors use as they share and assess each other’s knowledge (see Figure 2, #1). When Bill and his colleagues developed the modeling tool, they were able to establish a base common language that they could use to specify critical differences (i.e., size, shape, geometry, weight, functionality, etc.) and dependencies measured (i.e., drag coefficient and “skin” of the car). In the case of the B-150, the lexicon developed and embodied in the CFD tool served as a basis for jointly identifying what was consequential. A climate control engineer who worked on the B-150 summed this up nicely when he said, “We disagree sooner and know what we are disagreeing about more productively since we have a shared way to compare our design parameters.” In Figure 2, the characteristic of a common lexicon (#1) is placed at the bottom of the triangle to indicate its foundational role in supporting the other three characteristics. As indicated before, when novelty is present, a syntactic capacity remains necessary, but is not an adequate capacity to share and assess knowledge.

The second characteristic required to share and assess knowledge across a pragmatic boundary is that the actors involved need the ability to identify and learn about new differences and dependencies between them when novelty is present (See Figure 2, #2). For example, on the B-150 redesign team, each of the four groups had very different requirements and design preferences. The vehicle styling group wanted a new “sleek” design; the engine/power train group preferred a newly developed, higher horsepower engine; the climate control group wanted a larger grill to solve past problems; and the safety group wanted to see a bumper design no higher than 16.5 inches from the ground to minimize collision damage to other cars. With these design preferences specified, the dependencies across them also need to be identified and assessed. The CFD model represented these dependencies and identified their consequences

Figure 2 3-T Framework and the Four Characteristics of a “Pragmatic” Boundary Capability



through a simulation and measurement of the vehicle's "skin" and drag coefficient. For example, representing the consequences of a more powerful engine showed how it affected the slope of the hood, increased the vehicle's weight and grill size, and changed the location of the bumper in a concrete and rapid manner that would have been impossible using the clay model. The CFD tool provided a semantic capacity at the boundary where each group could identify the novelty present and translate their knowledge to develop common meaning and agreements about what is of consequence. When only a semantic boundary is faced, the groups involved do not have to transform their domain-specific knowledge; rather, the development of a common meaning is enough for them to effectively share and assess their knowledge.

The third characteristic required at a pragmatic boundary is a transformation of actors' domain-specific knowledge so they can work effectively together (see Figure 2, #3). Being able to propose, negotiate, and transform knowledge lies at the heart of trial-and-error problem solving at a pragmatic boundary. The groups involved must be able to "try on" alternatives and make trade-offs; new agreements can now be created to adequately share and assess knowledge at a boundary. Like the example of the firemen's tarp, the ability of actors to change their own and other's knowledge only emerges when there is a pragmatic capacity, a way of representing the consequences of how the knowledge of one group generates consequences on the knowledge of another group, and then making changes accordingly. By using the CFD modeling tool, each group could first represent their various concerns, data points, and requirements, then engage each other to identify, negotiate, transform, and verify the knowledge that they would then use to design the vehicle at this early stage. This assessment criterion is usually a collectively defined measurement

of cost, quality, time, or, in this case, the particular drag coefficient negotiated and the shape of "skin" eventually reached. Although the existence of a common measurement (as a lexicon) is required in using the tool, it is important to realize that the process of developing common meaning and interests cannot be made as explicit.

The fourth and last characteristic of managing knowledge at a pragmatic boundary is that it requires multiple iterations (see Figure 2, #4). Addressing the consequences cannot be resolved with one try, but requires an iterative process of sharing and assessing knowledge, creating new agreements, and making changes where needed. As the actors participate in each iterative stage, they get better at identifying what differences and dependencies are of consequence at the boundary; they improve at collectively developing more adequate common lexicon, meaning, and interests. Through this iterative capacity the invested and path-dependent nature of knowledge can be transformed. For example, on previous redesign efforts on the B-150, sufficient heating and cooling power of the climate control system had always been a significant source of expensive downstream design rework. The requirements of the climate control group had always been sandwiched between two powerful engineering groups: vehicle styling, which defined the size grill, and engine/power train, whose engine generates most of the heating and cooling requirements. However, because of the iterative capacity of the CFD tool to compare design requirements, identify consequences, experiment, and make trade-offs, the B-150 redesign was the first time that the grill (and its heating and cooling power) on the production vehicle followed the same overall specifications defined in the early design stage. The representational capacity of the CFD tool and the ability of the actors to use it helped reduce costs and delays on the B-150.

4.3. Understanding the Failure of the CFD Tool

Bill became concerned when downstream delays and costs rose on the B-100, despite the use of the CFD tool. At first he thought that the tool had not been used correctly because of training or resource limitations. A resource audit was conducted to assess staffing and training levels at the early design stage of the B-100, but no glaring problems were identified. Bill then turned his attention to the biggest source of the downstream costs and delays faced on B-100—is failure on three out of four primary crash tests. Given the increased safety consciousness of the public in the 1990s, crash-test reports had become a basic requirement in a competitive market. This would indeed be an expensive problem to fix—But what were the sources of this problem?

First, the B-100 is a smaller platform version of the B-150—The gross design and layout of major components were comparable. This fact was voiced by many members of the B-100 team along the following lines: “It was a similar vehicle, just on a smaller platform.” One assumed similarity that was not questioned thoroughly enough was how a “shrunk” design from the B-150 would perform on crash tests; therefore, some of the differences in the size and layout of the B-150 and the B-100 proved significant. Most critically, the crash impact of the bumper in relation to large structures of the vehicle, such as the frame, engine, and other large components, was not reconsidered for a smaller platform. In short, the “crash physics” and “trajectories of force” of the B-150 design generated negative impacts on the B-100 that were not identified in the early design stage.

Second, because of the technical similarity between the two vehicles, very early on in the use of the CFD tool the deployment engineer identified a drag coefficient of 0.33 as a good target because it had been one of the “successful” measures reached at the end of the B-150 redesign. Given the design similarities between the two platforms, the deployment engineer and others on the B-100 redesign team believed that “the target of 0.33 was technically a justifiable place to start for the B-100 redesign.” There were other specifications that were also similar to the B-150 that were easily identified and quickly agreed on; together, these specifications reinforced the assumed adequacy of 0.33 as a critical design constraint. The consequences of specifying this technical target at the beginning anchored and overly constrained further efforts to identify and learn about potential drawbacks of the current approach to the design. This proved particularly consequential because the drag coefficient was one of the two primary “measurements” of the dependencies and the potential consequences that each group placed on the other.

Third, although the safety engineer was experienced in vehicle design, he was new to the B-100 platform. There was a fair amount that he had to learn before he could

really determine what novelty he had to share and what novelty was important for him to assess to determine any negative consequence the current design had on safety issues. This lack of experience worked against the safety engineer in ways that he did not even realize because of the constrained nature of the process described above. Because a critical measure of dependence (0.33) was determined so early, many of the differences and dependencies that the safety engineer needed to represent to others were either not identified or, if identified, not sufficiently solved to identify all of the consequences that eventually proved so costly downstream (i.e., crash-test failures).

Even though the engineering tool used and the numeric outcomes reached were the same (i.e., 0.33 drag coefficient) across the B-150 and B-100, the process of getting there was not. On the B-100 the iterative process of representing, learning about what was most consequential, and transforming the design was constrained and truncated by the assumed similarities. The downstream costs that arose on the B-100 platform resulted from reusing a key piece of knowledge developed on another platform; conceptually the lexicon of 0.33 stood in the way of recognizing and resolving the lack of common meaning and interests on the B-100. More practically this limited the capacity and the ability of the actors—particularly the safety engineer—to represent and resolve the consequences of the novelty present. This set of features led to the “mismatch” at this pragmatic boundary and the significant downstream costs that resulted.

5. Discussion

The purpose of this paper has been to revisit the topic of boundaries through an examination of managing knowledge across specialized domains when innovation is a desired outcome. To address the varying circumstances that are possible at a boundary, the following relational properties of knowledge at a boundary were discussed: difference, dependence, and novelty. This provided a language to talk about the increasing effort required to adequately share and assess domain-specific knowledge as circumstances at a boundary grow more complex. The categories syntactic, semantic, and pragmatic were used to scale the increasing complexity of managing a boundary and were respectively linked to the information-processing, interpretive, and political perspectives in organization theory. The framework developed (Figure 1) helped to recognize and resolve many of the incompatibilities that these perspectives in organization theory often present.

Focusing on the effectiveness of managing knowledge across boundaries clarified the distinction between domain-specific knowledge and common knowledge at a boundary. Acknowledging the importance of matching

the capacity of the common knowledge (common lexicon, meaning, and interests) with the type of boundary faced, as well as making sure that the actors involved have the ability to use that common knowledge or expertise (Black et al. 2004), provided a “prescriptive” platform to address the challenges of managing knowledge across boundaries. The *capability* to manage knowledge at a boundary can be framed as a combination of capacity and ability—capacity times ability equals capability. Complicating this rough formula is the recognition that actors tend to reuse knowledge even when novelty is present. This invested, path-dependent tendency is the source of the competency traps and the “curse of knowledge” that generate mismatches between actors.

The conceptual and prescriptive value of the framework helps resolve such paradoxes as to why knowledge is often sticky within the silos of a firm but leaks to competitor firms (Brown and Duguid 2001). The framework easily outlines a number of empirical questions as to why this is the case: Across what type of boundary is knowledge sticky and at which is it leaky? What is the amount of difference, dependence, and novelty at each? What is the existing capability (capacity times ability) present at each boundary? Such a paradox could also be seen in the case. In the 1960s the clay model provided an adequate common knowledge to share and assess knowledge at the early stage of a vehicle’s design, but with the competitive changes in 1980s and 1990s, it no longer had the capacity to share and assess the knowledge that had to be identified and resolved. With the development of the CFD tool and its application to the B-150, we see a representational tool that had an adequate capacity to represent the relational properties of knowledge at this early stage and significantly reduce the delays and cost overruns. The CFD tool used on the B-100 platform still had sufficient representational capacity, but the team assumed that a critical measure of dependence (0.33) was an adequate starting point. As a common lexicon of dependence, it constrained the process the actors used to share and assess each other’s knowledge, limiting their ability to identify the novelty present and understand its consequences (i.e., crash-test physics on a smaller platform). This constrained the identification of the conflicts that would arise as well as the development of common meaning and interests to resolve them.

Understanding the different capacities of common knowledge (i.e., lexicon, meaning, and interests) and the abilities of the actors involved to use them improves our understanding of when an object can function as a boundary object (Star 1989). The distinction between types of boundaries reminds us that depending on the type of boundary faced, boundary objects with different capacities are required. For example, at an information-processing boundary where the circumstances are stable, a common lexicon (i.e., price, specifications, 0s and 1s) is a sufficient boundary object, whereas the lexicon

of 0.33 for the B-100 was not. As seen in the case discussed, when a semantic or pragmatic boundary is faced, the ability of the actors to use the boundary object can no longer be taken for granted. The majority of empirical research on boundary objects only describes their benefits (Bechky 2003, Henderson 1999, Ancona et al. 2001), not how and why they function the way they do. This framework explains why a given boundary object is no “magic bullet” (Carlile 2002) when it is used in a situation where its capacity as a type of common knowledge (lexicon, meaning, interests) and/or the ability of the actors to use it is not well matched.

A relational understanding of knowledge at a boundary also specifies at a very concrete level the relationship between knowledge and power (Foucault 1980, Hardy and Clegg 1996). Even when actors have equal ability to use a common knowledge to effectively share and assess each other’s domain-specific knowledge, power is still being expressed. In these circumstances the “relative” power of each actor to represent differences and dependencies to each other is roughly matched (Black et al. 2004); thus, the knowledge and power each actor uses does not generate negative consequences. However, when abilities to use the common knowledge are not equal or the common knowledge used does not have the capacity to represent a particular actor’s knowledge and interests, mismatches arise. Extending this, in complex processes in our society (i.e., product development, public policy development, etc.) specialized knowledge is distributed across different domains and cannot always be equally represented at the same time. This temporal dimension of dependency means that the consequences of downstream knowledge generally have a harder time being represented earlier in the process, putting upstream knowledge (i.e., designing a product or policy) in a politically stronger position relative to downstream knowledge (i.e., building the product or implementing the policy). Given this, we should not assume the actors involved at a boundary occupy politically equal positions in representing their knowledge to each other. This is why the clay model continued to be the dominant method used at this complex boundary even into the 1990s; it represented what was at stake for the most upstream and powerful group, vehicle styling and their marketing champion. In the 1990s safety was the newest group and was in the weakest position to represent its domain-specific knowledge. This is why the effectiveness of an outcome at a pragmatic boundary is based on the capacity of the common knowledge and the ability of the actors involved to use it.

As with any research there are limitations in this work; three in particular require specific mention. First, the case itself, although illustrative, does not in any way test the framework being developed. However it should be pointed out that the development of this framework arises out of several empirical studies on sharing

and managing knowledge across specialized domains (Carlile 2002, Carlile and Reberntsch 2003, Black et al. 2004); thus, the major concepts in the framework have been empirically and conceptually refined. Second, this work emphasized the desirability of convergence rather than divergence at a boundary. Given the time horizons and performance constraints that product-development organizations face, convergence is a desired outcome. However, in more loosely coupled systems such as science and the arts, a certain amount of divergence generates more robust outcomes in the long term (i.e., physics; see Galison 1999). Last, as with any work that attempts to conceptualize and incorporate insights from across different disciplinary domains, the effort is not without its rough edges. To integrate different perspectives, the framework simplifies some issues while stressing others. This effort can be seen as a strength or weakness, depending on how broad an approach to boundaries one takes. The irony of this is that this work is itself an example of how hard it is to work across boundaries—to produce a common knowledge that can be of value in a multidisciplinary field such as organization theory.

6. Implications

The fact that most innovation occurs at the boundaries between specialized domains (Leonard-Barton 1995) tells us that effectively managing knowledge across the various types of boundaries in an organization is what drives competitive advantage. Applying this framework to strategic questions provides a concrete way to describe core concepts such as dynamic capability (Teece et al. 1997), where the stated concern has been how to change old knowledge to create new knowledge in a firm. The case described the capability of this process at a pragmatic boundary and provided an example of what is required (or lacking). Failures occurred when the actors involved did not have sufficient capacity or ability to manage the novelty that was present. At a firm level, a dynamic capability can be thought of as a collection of different combinations of capacities and abilities that can be used to share and assess knowledge across the various types of boundaries. From this vantage point, instead of seeing the firm as a bundle of resources (Barney 1991), it can be more completely described as a bundle of different types of boundaries where knowledge must be shared and assessed. This framework provides a concrete description of such boundary capabilities (“capacity times ability”), a potentially fruitful ground to begin linking a firm’s organizational and strategic views. This work is also an addition to the knowledge-based view of the firm (Grant 1996, Spender 1996) that has also been suggested as a connection between organizational and strategic views. The benefits of this work is that it provides conceptual specificity to knowledge as both a thing (content) and a process, and describes the

consequences of path dependency and power that have not been addressed by knowledge-based views.

More broadly, the potential conceptual value of the framework is its application to a number of classic and more contemporary topics such as differentiation and integration (Lawrence and Lorsch 1967), boundary spanning (Allen 1977), organization design (Galbraith 1973, Tushman and Nadler 1978), absorptive capacity (Cohen and Levinthal 1990), and modularity (Baldwin and Clark 1999). For each, the question of boundary management when novelty arises must be addressed; however, what is generally left unspecified by these topics is what type of boundaries are being faced, what is required to develop an adequate common knowledge, and how the current capacities and abilities might need to be changed to address the novelty now present. The challenge for any perspective that seeks to explain how systems adapt over time remains understanding the capacities and the abilities of actors to make the necessary trade-offs between the knowledge that was used before and the novelty present to create something new.

This has implications for a broad definition of actors, whether they are individuals, groups, organizations, or even computer-based agents. This challenge, however, is magnified by the tendency of actors to reuse knowledge, which limits their capacity and ability to represent differences and dependencies when novelty is present. This position of power—at least in the short term—whether exercised consciously or not, makes it harder for novelty to be represented and the consequences of it understood. These practical and political challenges need to be better recognized to explain how to generate innovation when knowledge must be shared and assessed by actors across different domains.

Acknowledgments

First, the author would like to thank Sim Sitkin and Ann Clark at *Organization Science* and the two anonymous reviewers for their expert help. Second, he would like to thank his colleagues at MIT who helped him improve the ideas in this paper: Lotte Bailyn, Rebecca Henderson, Wanda Orlikowski, Eric Reberntsch, John Van Maanen, and Eleanor Westney. Last, he would like to thank MIT’s Center for Innovation in Production Development for help in funding this research.

Endnote

¹“Common knowledge” is used in the more technical sense (i.e., communication theory) of referring to a shared body of knowledge that allows for communication between actors (see also “mutual knowledge,” Cramton 2001).

References

- Allen, T. 1977. *Managing the Flow of Technology*. MIT Press, Cambridge, MA.
- Ancona, D., D. Caldwell. 1992. Bridging the boundary: External activity and performance in organizational teams. *Admin. Sci. Quart.* 37 634–665.

- Ancona, D., G. Okhuyson, L. Perlow. 2001. Time-out: Taking time to integrate temporal research. *Acad. Management Rev.* **26** 512–529.
- Argote, L. 1999. *Organizational Learning: Creating, Retaining, and Transferring Knowledge*. Kluwer, Norwell, MA.
- Arrow, K. 1994. Methodological individualism and social knowledge. *Amer. Econom. Rev.* **84**(2) 1–9.
- Baldwin, C., K. Clark. 1999. *Design Rules, Vol. 1: The Power of Modularity*. MIT Press, Cambridge, MA.
- Barley, S. R. 1986. Technology as an occasion for structuring: Evidence from observations of CT scanners and the social order of radiology departments. *Admin. Sci. Quart.* **31** 78–108.
- Barney, J. B. 1991. Firm resources and sustained competitive advantage. *J. Management* **17** 99–120.
- Bechky, B. 2003. Sharing meaning across occupational communities: The transformation of understanding on the production floor. *Organ. Sci.* **14** 312–330.
- Black, L., P. R. Carlile, N. Repenning. 2004. Expanding theoretical insights from ethnographic evidence: Building on Barley's study of CT-scanning implementations. Pending publication.
- Bourdieu, P. 1980. *The Logic of Practice*. Cambridge University Press, Cambridge, U.K.
- Bourdieu, P., L. Wacquant. 1992. *An Invitation to Reflexive Sociology*. University of Chicago Press, Chicago, IL.
- Brown, J. S., P. Duguid. 1991. Organizational learning and communities-of-practice. *Organ. Sci.* **2** 40–57.
- Brown, J. S., P. Duguid. 2001. Knowledge and organization: A social-practice perspective. *Organ. Sci.* **12** 198–213.
- Brown, S., K. Eisenhardt. 1995. Product development: Past research, present findings, and future directions. *Acad. Management Rev.* **20** 343–378.
- Camerer, C., G. Lowenstein, M. Weber. 1989. The curse of knowledge in economic settings: An experimental analysis. *J. Political Econom.* **97** 1232–1254.
- Carlile, P. 2002. A pragmatic view of knowledge and boundaries: Boundary objects in new product development. *Organ. Sci.* **13** 442–455.
- Carlile, P., W. Lucas. 2003. Taking care of complex boundaries: Knowledge and boundary activities on technology development teams. Working paper, Sloan, MIT, Cambridge, MA.
- Carlile, P., E. Rebentisch. 2003. Into the black box: The knowledge transformation cycle. *Management Sci.* **49** 1180–1195.
- Christensen, C. M. 1997. *The Innovator's Dilemma: When New Technologies Cause Great Firms to Fail*. Harvard Business School Press, Boston, MA.
- Cohen, W., D. Levinthal. 1990. Absorptive capacity: A new perspective on learning and innovation. *Admin. Sci. Quart.* **35** 128–152.
- Cramton, C. 2001. The mutual knowledge problem. *Organ. Sci.* **12** 346–371.
- Crowston, K. 1997. A coordination theory approach to organizational process design. *Organ. Sci.* **8** 157–175.
- Cruse, D. 2000. *Meaning in Language: An Introduction to Semantics and Pragmatics*. Oxford University Press, New York.
- Davenport, T., L. Prusak. 1998. *Working Knowledge*. Harvard Business School Press, Boston, MA.
- Dougherty, D. 1992. Interpretive barriers to successful product innovation in large firms. *Organ. Sci.* **3** 179–202.
- Foucault, M. 1980. *Power/Knowledge: Selected Interviews and Other Writings 1972–1977*. Harvester Press, Brighton, Sussex, U.K.
- Galbraith, J. 1973. *Designing Complex Organizations*. Addison-Wesley, Reading, MA.
- Galison, P. 1999. *Images and Logic: A Material Culture of Micro-Physics*. University of Chicago Press, Chicago, IL.
- Gittell, J. 2001. Supervisory span, relational coordination and flight departure performance: A reassessment of postbureaucracy theory. *Organ. Sci.* **12** 468–483.
- Grant, R. 1996. Toward a knowledge-based theory of the firm. *Strategic Management J.* **17** 109–122.
- Hardy, S., S. Clegg. 1996. Some dare call it power. S. Clegg, C. Hardy, C. Nord, eds. *Handbook of Organization Studies*. Sage, London, U.K.
- Hargadon, A., R. Sutton. 1997. Technology brokering and innovation in a product development firm. *Admin. Sci. Quart.* **42** 716–749.
- Henderson, K. 1999. *On Line and on Paper: Visual Representations, Visual Culture, and Computer Graphics in Design Engineering*. MIT Press, Cambridge, MA.
- Henderson, R., K. Clark. 1990. Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Admin. Sci. Quart.* **44** 83–111.
- Hinds, P. 1999. The curse of expertise: The effects of expertise and de-biasing methods on predictions of novice performance. *J. Experimental Psych.* **5** 205–221.
- James, W. 1907. *Pragmatism*. The American Library, New York.
- Lave, J. 1988. *Cognition in Practice: Mind, Mathematics, and Culture in Everyday Life*. Cambridge University Press, New York.
- Lave, J., E. Wenger. 1991. *Situated Learning*. Cambridge University Press, Cambridge, U.K.
- Lawrence, P., J. Lorsch. 1967. *Organizations and Environments: Managing Differentiation and Integration*. Harvard Business School Press, Cambridge, MA.
- Levitt, B., J. G. March. 1988. Organizational learning. W. R. Scott, J. Blake, eds. *Annual Review of Sociology*. Annual Reviews, Palo Alto, CA, 319–340.
- Leonard-Barton, D. 1992. Core capabilities and core rigidities: A paradox in managing new product development. *Strategic Management J.* **13** 111–126.
- Leonard-Barton, D. 1995. *Well Springs of Knowledge: Building and Sustaining the Sources of Innovation*. Harvard Business School Press, Boston, MA.
- Litwak, E., L. F. Hylton. 1962. Interorganizational analysis: A hypothesis on coordination agencies. *Admin. Sci. Quart.* **6** 395–420.
- Malone, T., K. Crowston. 1994. The interdisciplinary study of coordination. *ACM Comput. Surveys* **26**(March) 87–119.
- March, J. 1972. Model bias in social action. *Rev. Ed. Res.* **44** 413–429.
- Martins, L., A. Kambil. 1999. Looking back and thinking ahead: Effects of prior success on managers' interpretations of new information technologies. *Acad. Management J.* **42** 652–661.
- Nonaka, I. 1994. A dynamic theory of organizational knowledge creation. *Organ. Sci.* **5** 14–37.
- Nonaka, I., I. Takeuchi. 1995. *The Knowledge-Creating Organization*. Oxford Press, Oxford, U.K.
- Orr, J. 1996. *Talking About Machine*. Cornell University Press, Ithaca, NY.
- Perrow, C. 1994. Accidents in high risk systems. *Tech. Stud.* **1** 1–38.

- Polyani, M. 1966. *The Tacit Dimension*. Anchor Day Books, New York.
- Reddy, M. J. 1979. The conduit metaphor. A. Ortony, ed. *Metaphor and Thought*. Cambridge University Press, Cambridge, U.K.
- Schank, R. C., R. Abelson. 1977. *Scripts, Plans, and Goals*. Erlbaum, Hillsdale, NJ.
- Schrage, M. 1999. *Serious Play: How the World's Best Companies Simulate to Innovate*. Harvard Business School Press, Boston, MA.
- Shannon, C., W. Weaver. 1949. *The Mathematical Theory of Communications*. University of Illinois Press, Urbana, IL.
- Spender, J. C. 1996. Making knowledge the basis of a dynamic theory of the firm. *Strategic Management J.* **17** 45–62.
- Star, S. L. 1989. The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving. M. Huhns, L. Gasser, eds. *Readings in Distributed Artificial Intelligence*. Morgan Kaufman, Menlo Park, CA.
- Szulanski, G. 1996. Exploring external stickiness: Impediments to the transfer of best practice within the firm. *Strategic Management J.* **17** 27–43.
- Teece, D. J., G. Pisano. A. Shuen. 1997. Dynamic capability and strategic management. *Strategic Management J.* **18** 509–533.
- Thompson, J. 1967. *Organizations in Action*. McGraw-Hill, New York.
- Tsoukas, H. 2001. Re-viewing organisation. *Human Relations* **54** 7–12.
- Tushman, M., D. Nadler. 1978. Information processing as an integrating concept in organization design. *Acad. Management Rev.* **3** 613–624.
- Tyre, M., E. von Hippel. 1997. The situated nature of adaptive learning in organizations. *Organ. Sci.* **8** 71–83.
- Victor, B., R. S. Blackburn. 1987. Interdependence: An alternative conceptualization. *Acad. Management Rev.* **12** 486–498.
- von Hippel, E. 1994. The impact of sticky data on innovation and problem solving. *Management Sci.* **40** 429–439.
- Weber, M. 1924/1947. *The Theory of Social and Economic Organization*. A. H. Henderson, T. Parsons, eds. Free Press, Glencoe, IL.
- Weick, K., K. Sutcliffe, D. Obstfeld. 1999. Organizing for high reliability: Processes of collective mindfulness. B. Staw, R. Sutton, eds. *Research in Organizational Behavior*. JAI Press Inc., Stamford, CT, 81–117.
- Wenger, E. 1998. *Communities of Practice: Learning, Meaning and Identity*. Cambridge University Press, Cambridge, U.K.
- Wheelwright, S., K. Clark. 1995. *Leading Product Development: The Senior Manager's Guide to Creating and Shaping the Enterprise*. Free Press, New York.
- Winter, S. 1987. Knowledge and competence as strategic assets. D. Teece, ed. *The Competitive Challenge*. Ballinger, Cambridge, MA.