

A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development*

Paul R. Carlile

Sloan School of Management, Massachusetts Institute of Technology, E52-567,
Cambridge, Massachusetts 02142-1347
carlile@mit.edu

Abstract

This study explores the premise that knowledge in new product development proves both a barrier to and a source of innovation. To understand the problematic nature of knowledge and the boundaries that result, an ethnographic study was used to understand how knowledge is structured differently across the four primary functions that are dependent on each other in the creation and production of a high-volume product. A pragmatic view of “knowledge in practice” is developed, describing knowledge as localized, embedded, and invested *within* a function and how, when working *across* functions, consequences often arise that generate problematic knowledge boundaries. The use of a boundary object is then described as a means of representing, learning about, and transforming knowledge to resolve the consequences that exist at a given boundary. Finally, this pragmatic view of knowledge and boundaries is proposed as a framework to revisit the differentiation and integration of knowledge.

(*Knowledge; Knowledge Management; Boundary Objects; Ethnography; New Product Development*)

Introduction

The topic of knowledge in organizations has had a great deal of attention in the literature for nearly a decade. For both researchers and practitioners, much of this intense interest is driven by the recognition that knowledge is a critical factor in creating competitive success over time (Kogut and Zander 1992, Nonaka and Takeuchi 1995). Many have realized, however, that executing on this

awareness is a significant challenge. The difficulty of transferring knowledge (Suzlanski 1996), the tacit nature of knowledge (Polanyi 1966, Nonaka 1994, von Krogh et al. 2000), and its stickiness (von Hippel and Tyre 1996) have been important revelations into why knowledge is difficult to “manage.” This paper continues in this vein, but explores this difficulty at a more concrete level to explain why knowledge remains a critical but challenging source of competitive advantage for an organization.

I start with the premise that knowledge in organizations is problematic; specifically, in new product development, knowledge is both a source of and a barrier to innovation. The characteristics of knowledge that drive innovative problem solving *within* a function actually hinder problem solving and knowledge creation *across* functions. It is at these “knowledge boundaries” that we find the deep problems that specialized knowledge poses to organizations. The irony is that these knowledge boundaries are not only a critical challenge, but also a perpetual necessity because much of what organizations produce has a foundation in the specialization of different kinds of knowledge.

To better understand these knowledge boundaries, I conducted a year-long ethnographic study to examine how knowledge is structured differently *within* the four primary functions (sales/marketing, design engineering, manufacturing engineering, and production) involved in the creation and production of a high-volume product. From this fieldwork I describe knowledge as *localized*, *embedded*, and *invested* in practice (Bourdieu 1977, Lave 1988). This specialization of “knowledge in practice” (Carlile 1997) makes working *across* functional boundaries and accommodating the knowledge developed in another practice especially difficult. The following interaction between a design engineer and a manufacturing engineer is an example of the knowledge boundaries that are generated as each

*Note. This paper was accepted for the *Organization Science* “Special Issue on Knowledge, Knowing, and Organizations.” Administrative constraints on issue size and timing dictated that the paper be published independently in the current issue.—Claudia Bird Schoonhoven, Editor-in-Chief

individual problem solves different objects toward the ends required within their respective practices.

Mick had been working on the on-board vapor recovery valve (OVRV) project for about four months now. He had already been to several design review meetings, and was beginning to anticipate a lot of problems in trying to assemble and test a new product with so many parts and critical sealing surfaces. As the sole representative from manufacturing engineering at this early design phase, he could easily summarize his frustration to me: “they [the design engineers] don’t realize that the OVRV, with its high part count and 3,000,000-a-year volume, is going to be a completely different beast to deal with.” Mick had already strongly expressed the necessity of going to subassemblies to make assembly and testing easier, but so far he hadn’t gotten any significant design changes approved. At each successive design meeting, the only critical changes he noticed were ones that improved the valve’s “functionality.”

In this design review meeting, Mick’s statements about the “awkwardness of the design” only seemed to be inflaming the tempers of the sales representative and design engineers on the team. For Vaughn, the head design engineer, the past challenges of packing all of the functional requirements of such a complex “3-in-1” valve into such a small space made the success of the “working prototype” and the current design a significant achievement. However, what increased the tension in the room for everyone was the announcement that the customer was holding firm on having test parts delivered in eight weeks. Mick kept up with his arguments about subassemblies, and he would later say to me, “what am I going to do, sit and wait for the production launch (of the OVRV) for all hell to break loose? I’ll have more than egg on my face then.”

In what follows, I first summarize two current approaches of how boundaries have been conceptualized in the product development literature and then propose a third. Second, I present my theoretical and empirical arguments of knowledge as localized, embedded, and invested in the particular objects and ends of a given function to frame the consequences of moving knowledge across boundaries. Third, I describe the use of “boundary objects” (Star 1989, Carlile 1997) by individuals and the characteristics of effective ones as they provide a means of resolving the consequences that arise when different kinds of knowledge are dependent on each other. Finally, I offer up this pragmatic view to reframe the classic model of differentiation and integration (Lawrence and Lorsch 1967) when applied to knowledge in settings where innovation across different functional specialties is a required outcome. This reframing suggests that “adequate capacity” (Galbraith 1973) is not just one of processing or transferring more knowledge, but “transforming knowledge” (Carlile 1997) to effectively deal with differences, dependencies, and the novelty present at a given boundary.

Three Approaches to “Knowledge Boundaries” in Product Development: Syntactic, Semantic, and Pragmatic

In what follows I describe three different approaches (Jantsch 1980) to knowledge and moving it across boundaries in the product development literature. The first two, *syntactic* and *semantic*, capture the current ways of thinking about knowledge and boundaries. I propose a third, *pragmatic*, as additional but complementary to the first two approaches. Together each approach successively frames increasingly complex and challenging boundaries that exist in new product development.

Syntactic Approach

A *syntactic* approach to boundaries was first developed by Shannon and Weaver (1949) with their establishment of a mathematical theory of communication. Once a syntax (e.g., a hexadecimal zero or one as in software code) is shared and stable across a given boundary, it is sufficient and the processing of information becomes the primary concern. For Shannon and Weaver, establishing a shared and stable syntax meant that they could ensure accurate communication between sender and receiver across a boundary and solve many challenging communication (i.e., avoiding friendly fire in combat) or information processing problems. This information processing perspective of dealing with a boundary also had significant impact on the social sciences. Perhaps most dramatic was with the “systems” theorists (e.g., Bertalanffy 1956, Ashby 1956, Buckley 1968) who framed the boundary between an organization and its environment as a problem to be solved by information processing.

Influenced by this systems approach, Lawrence and Lorsch introduced their differentiation and integration model (1967) to frame the challenge an organization faces in effectively dealing with its environment. Differentiation arises because subunits (i.e., sales, R&D, production) face environments with different degrees of uncertainty. Lawrence and Lorsch measured this source of difference according to the degree of task predictability in each specialized subunit (Grandori 1986). This common interval measure allowed them to sufficiently describe differentiation across all subunits of the organization, which narrowed the problem of integration to one of “matching differences” (Lawrence and Lorsch 1969, p. 26) or degrees of uncertainty on either side of the boundary. It is through the existence of a shared and stable syntax across a boundary that matching occurs and insures a “quality information exchange” (1967, p. 33). This approach, coupled with Galbraith’s (1973) requirement of establishing “adequate information processing

capacity,” made the information processing model and its syntactical requirements the dominant approach to boundary spanning in organization theory and new product development research (Brown and Eisenhardt 1995).

The first to apply Lawrence and Lorsch’s (1967) approach within a product development setting was Allen (1971, 1977) in his work examining the communication patterns and the distances between individuals involved in product development efforts and their impact on success. Many extended this boundary spanning and information processing framework and focused on what internal communication patterns, planning, and prioritizing processes were determinants of product success (Keller 1986, Joyce 1986, Ancona and Caldwell 1992a). Others focused on the importance of external communication patterns and boundary spanning activities in successful product development (Katz and Tushman 1981, von Hippel 1988, Ancona and Caldwell 1992b). The overall insight from this type of research is that more information is better, more communication is better, and more team strategies are better. However, beyond this basic insight, a more critical issue remains: When novel conditions emerge will the current syntax be sufficient to process information at the boundary? Given this, the resource problem shifts from one of processing more information to understanding these novel conditions or new knowledge that lies outside the current syntax used at the boundary.

In the example presented in the introduction, the high part count of the OVRV and its volume requirement of 3,000,000 a year was significantly higher than any previous product manufactured. This presented Mick with a different and novel problem to be solved. This new requirement made the old syntax (i.e., no subassemblies) insufficient to begin problem solving the boundary between Mick and Vaughn, their new circumstances or task that required “subassemblies” demanded new and different syntax to support effective communication between them. This example shows that novelty does not arise from differences in degree of uncertainty but, more problematically, differences in kinds of knowledge required for the task. To describe the circumstances when a syntactical approach is no longer sufficient, we turn to a perspective that frames novel sources of difference across a boundary.

Semantic Approach

A *semantic* approach recognizes that even if a common syntax or language is present, interpretations are often different which make communication and collaboration difficult. The interpretive or cultural perspective in social science over the last 30 years has brought to the fore the

importance of recognizing that interpretive differences exist and that messages are often problematic (Redding 1972, Reddy 1979). A semantic approach also recognizes that difference is not always adequately represented as “differences in degree” (i.e., differentiation), but “differences in kind.” The problem then shifts from just processing information to learning about the sources that create these semantic differences that exist at a boundary.

Compared to the traditional syntactic approach, a semantic approach to boundary spanning in product development is a much smaller literature. Dougherty’s work (1992) on “thought worlds” (Fleck 1935) offers a nice example of a semantic perspective on why differences in meaning or language across functions in product development remain challenging. In her work, she outlines how these different thought worlds (what they know and how they know it) make communications difficult because individuals use different meanings in their functional setting.

Nonaka and Takeuchi (1995, p. 58) recognize the contrast between the syntactical and the semantic approach when they clarify the particular challenges in trying to create and transfer knowledge in organizations. By paying attention to the challenges of “conveyed meaning” and the possible different interpretations by individuals, they recognize that individual, context-specific aspects of creating and transferring knowledge must be taken into consideration, and like others pay particular attention to the tacit nature of knowledge (Polanyi 1966, Leonard-Barton 1995, von Hippel and Tyre 1996). Nonaka (1994) suggests that what is required is the generation of “mutual understanding” through communities of interaction where individuals can work through these semantic differences by making tacit knowledge explicit across a boundary. Such a process focuses on the practical process of learning about and making explicit new sources of difference (e.g., Mick’s 3,000,000-a-year volume and Vaughn’s high number of critical sealing surfaces in a very small space). However, learning about differences is not always enough to deal with every knowledge boundary; in some cases by making one’s knowledge explicit the potential conflicts and costs associated in working across a boundary are made more explicit.

The semantic approach embraces sources of difference, but while recognizing dependency it does not acknowledge the consequences that are often generated because of such dependence. We understand the problematic nature of knowledge when we see not only how different kinds of knowledge create novelty (i.e., Vaughn’s “successful” current design and Mick’s subassembly solution to high part count, high scrape rate concerns), but also

how dependence across these differences generate consequences as well.

Pragmatic Approach

The pragmatic approach, with its roots in the philosophies of Peirce (1898) and James (1907), highlights the importance of understanding the consequences that exist between things that are different and dependent on each other. Bourdieu's (1977) "within"-practice empirical focus frames how knowledge is geared to make a particular effect (e.g., solve a particular problem), and because of that individuals are committed to and invested in their knowledge as hard-won outcome (i.e., for Mick to build a 3,000,000-a-year manufacturing process; for Vaughn to deliver a "3-in-1" valve that meets functional and size requirements). Further, Bourdieu's relational structuralism (Bourdieu and Wacquant 1992) also describes that interactions across practices are not inconsequential; the knowledge that people accumulate and use is often "at stake." They are reluctant to change their hard-won outcomes because it is costly to change their knowledge and skills. The cross-boundary challenge is not just that communication is hard, but that to resolve the negative consequences by the individuals from each function they have to be willing to alter their own knowledge, but also be capable of influencing or transforming the knowledge used by the other function.

In the case of Mick, the new knowledge that Vaughn developed to meet his requirement of a "3-in-1" valve in a small space was creating significant consequences for Mick's ability to develop a solution to his 3,000,000-a-year volume requirement. For Mick, his proposal to deal with those consequences of a significantly more complex valve was to go to a subassembly design. As this was a novel and unknown approach to the others on the design team, the current design would need to be altered in ways that could negatively impact the current performance of the OVRV design.

A pragmatic approach to boundaries assumes the conditions of difference, dependence and novelty are all present, and so recognizes the requirement of an overall process for transforming existing knowledge to deal with the negative consequences that arise. Here, transforming knowledge (Carlile 1997) refers to a process of altering current knowledge creating new knowledge, and validating it within each function and collectively across functions. For example, Schrage's (1999) work on "serious play" in driving innovation and Iansiti's (2000) work on the role of prototypes are observations that are suggestive of a pragmatic process of transforming knowledge.

A Pragmatic View of Knowledge: Localized, Embedded, and Invested in Practice

The community of practice literature has been particularly helpful in looking at how knowledge and learning is structured by the types of problems faced within a practice. This work was developed by Lave and Wenger (1991), Brown and Duguid (1991), and Orr (1996); it recognized the situated and "purposive" nature of knowledge as it is created by a community of individuals who have a shared practice or problem and share in its consequences. This work has clarified the situated and tacit characteristics of knowledge (Suchman 1987, Cook and Brown 1999) to remind us that it is not enough to have a shared syntax to work across communities of practice. This insight suggests that knowledge is not only tacit, in the narrow sense of it being that which is not explicit (Nonaka 1994), but also that knowledge and knowing cannot be separated from an individual's engagement in the "practicing" of their practice (Cook and Brown 1999).

Building on this research, this paper purposefully casts a broader empirical net to examine knowledge across four communities of practice involved in a new product development setting (sales/marketing, design engineering, manufacturing engineering, and production) and to examine how different kinds of knowledge impact working "across" communities of practice. The community of practice literature has previously focused empirically on a single community of practice, which has led it to take an interpretive or a semantic approach to the challenges of working across practices. However, the ethnographic approach taken here compares and contrasts the knowledge used across four practices (Carlile 1997) and clarifies three characteristics of knowledge in practice—that knowledge is localized, embedded, and invested in practice.

First, knowledge in new product development is localized around particular problems faced in a given practice. The effective development of knowledge in organizations demands that individuals specialize or localize around different problems. To say localized does not mean that knowledge is limited to only one situation or location; rather, knowledge can be quite similar across practices if it is localized around a similar set of problems; knowledge is local in character, not global. For example, systems engineering is a method used in dealing with large-scale engineering projects. To be useful, systems engineering methodologies have to "black box" many other types of problems faced inside the system. As this example suggests, even if a large-scale problem exists, its solution does not offer up a "global" (Wenger 1998,

p. 140) account of all of the knowledge required to successfully implement a global solution.

Second, knowledge is embedded in practice. The word “embedded” suggests an archaeological image as to why knowledge is hard to articulate or recall, knowledge accumulated in the experiences (Taylor 1992) and know-how (Harper 1987) of individuals engaged in a given practice. Knowledge is also embedded in the technologies, methods, and rules of thumb used by individuals in a given practice. Emphasizing that knowledge is embedded in practice gets us closer to examining the sources of why “we know more than we can tell,” which was stressed by Polanyi (1966) and his emphasis on “tacit knowing” residing in the doing of the activity. Seeing knowledge as embedded in practice offers a contrast to the cognitive expression of tacit knowledge as something that is hard to retrieve from the mind that defines much of the literature. The bottom line is that the more “distance” individuals have from each other’s practice—their engagement in practice—the more difficult it is to communicate the embedded knowledge they use.

Third, knowledge is invested in practice—invested in the methods, ways of doing things, and successes that demonstrate the value of the knowledge developed. When knowledge proves successful, individuals are inclined to use that knowledge to solve problems in the future. In this way, individuals are less able and willing to change their knowledge to accommodate the knowledge developed by another group that they are dependent on. Changing their knowledge means an individual will have to face the costs of altering what they do to develop new ways of dealing with the problems they face. Knowledge is one of the means by which individuals demonstrates their competency in solving problems to others inside (I’m a good design engineer) and outside (I met my requirements) their practice. This is why deliverables and deadlines are “at stake” and are for the individual a measure of their competency and success. The harsh reality is that the current knowledge accumulated to achieve these deliverables and deadlines will be “at stake” in future problem-solving efforts at a boundary.

These three characteristics of knowledge in practice have significant positive benefits *within* a practice. For knowledge to be useful and effective in solving the problems, individuals must be able to localize knowledge around particular problems, as well as draw from and alter (i.e., trial and error, and learning) the knowledge embedded within their practice. Individuals also must be invested in their knowledge as they try to meet the challenging requirements (what is “at stake”) in their practice. However, these same characteristics of knowledge in

practice that lead to the effective specialization of knowledge become problematic when working *across* practices.

Empirical Lens: Observing Objects and Ends in Practice

To be able to empirically describe and then theorize about knowledge boundaries, I developed a simple framework to compare how knowledge is structured in practice by focusing on the objects and ends used in a given practice. “Objects” refer to the collection of artifacts that individuals work with—the numbers, blueprints, faxes, parts, tools, and machines that individuals create, measure, or manipulate. “Ends” are outcomes that demonstrate success in creating, measuring, or manipulating objects—a signed sales contract, ordering prototype parts, an assembly process certification, or a batch of high-quality parts off the production line. The work itself is an ongoing process of moving an object from its current state to a required end state. These iterative problem-solving activities should be seen broadly as encompassing the know-how, techniques, and “trial and error” that it takes to move one’s objects toward a required end.

I took an anthropological approach (Spradley 1979, 1980; Reynolds 1987) that focuses on how individuals in each function apply their knowledge to the objects and ends that are of consequence in their practice. This approach was taken for two reasons. First, collecting data about knowledge has proven difficult. Given my interest in knowledge in practice, surveys and relying on interviews is not a direct enough method of getting at knowledge. Observing individuals in practice and focusing on the objects they work with and the ends that they pursue provides a concrete delineation of what to observe and what to compare in terms of how knowledge is created and structured. Second, the focus on the different objects and ends would allow for an interesting comparison, and would also provide specific examples of “objects and ends” that are dependent on each other in developing a successful product. This provided data about why knowledge is “at stake” within a practice, and how this makes the consequences between the differences and dependencies across practices difficult to resolve.

Collecting and Representing Data About Knowledge and Knowledge Boundaries

The observations that follow about knowledge in practice were collected over a period of a year (June 1994 to May 1995) using a field-based methodology. The purpose of this fieldwork was two-fold. The first is to understand how practice shapes knowledge and the problematic boundaries that exist between different functions in product development, the second to identify what activities or

processes were effective in facilitating collaboration across knowledge boundaries. In a practice-based research approach, it is crucial to be able to observe what people do, what their work is like, and what effort it takes to problem solve their respective combinations of objects and ends. Also, in comparing multiple practices, informants seldom have the experiential opportunity to make useful comparative observations about different practices. Because of this I chose a small firm (300 people, all in one site) that made it easier to observe most of what went on within each function and across those functions.

Field Site. X.T. Products (a pseudonym) designs and manufactures safety and environmental valves for automobile fuel systems and supplies them to both domestic and foreign automobile manufacturers. The majority of X.T.'s valves are "rollover," "pressure relief," and "cut" valves. In the case of a rollover accident, these valves cut off the flow of gas to the engine compartment to reduce the likelihood of fire. I focused on the product development process of a new, more complex valve called the OVRV, or "On-board Vapor Recovery Valve." The OVRV would be three valves in one: It would be a rollover valve, a pressure relief valve, and a vapor recovery valve. The vapor recovery valve was designed to meet new environmental requirements of minimal gas vapor loss into the atmosphere when pressure is relieved or when filling the gas tank. Integrating three valves into one while meeting the government standard meant the product would have nearly 60 parts, well over four times the number of parts in any previous valve. Though it would be much larger than previous valve, its design would have to pack more functionality into a comparatively small space.

I committed three to four days a week onsite for nearly a year and observed individuals and groups (i.e., watched, listened to, talked with, or questioned) as they worked and as I talked with them about their work. I also observed individuals as they worked in cross-functional settings. In total, I documented 106 cross-functional events with team members, both in formal meetings and informally around a desk or a "machine" as conversations or issues arose.

Representing Observations: Two Sets of Vignettes. The first set of data is a series of four short vignettes that depict knowledge in practice across sales, design engineering, manufacturing engineering, and production. Table 1 provides a summary of the objects and ends in each practice from a larger set of ethnographic work (Carlile 1997) that this paper draws from. These vignettes provide the reader with a set of comparisons of how knowledge is localized, embedded, and invested in solving different

problems. The second set of data is a larger vignette that continues the story from the introduction of Mick and his "frustration" in attempting to get Vaughn and others to change the design of the OVRV to accommodate subassemblies. The focal point of this vignette is the use of a "boundary object" to transform the current design of the OVRV to resolve the negative consequences given Mick and Vaughn's differences and dependencies. A challenge that every ethnographer faces is how to represent and make understandable a large collection of observations of day-to-day work or practice. In this research, I include these vignettes to represent the veracity of the knowledge that individuals use, as descriptive of what I saw unfold. They provide a representative sample of the knowledge in practice from the larger ethnographic study (Carlile 1997).

Vignettes of Knowledge in Practice in Different Functional Settings

Sales Work: Getting the Numbers "Right"

Ken is in charge of sales for the OVRV. This story focuses on how the numbers (i.e., volume requirements, specification, costs, etc.) are not only an identifiable outcome in his negotiations with the customer and upper management but, more dramatically, are concrete "stakes" in the ground that measures Ken's hard-won success.

As far as the day-to-day grind, the biggest struggle for Ken was that he was not only dealing with the customer, but his own upper management as well. Although upper management was vitally committed to getting this business, they were not about to "give too much away to get it." This competitive struggle, with Ken sitting in between the needs of the customer and the needs of upper management, resulted in nearly 11 weeks of intense negotiation: formally, informally, person-to-person, over the phone, and via the fax. The numbers to be negotiated included product prices, delivery schedules, and product specifications.

The final version of these negotiations, the numbers and specifications, go into what is called the "blue book." Ken remarked how small the final copy looks, and how its "tidiness is an insulting disguise to all of the late nights, pain, and frustration that went into getting there." But at the same time, he expressed a sense of relief in getting the numbers "right," so that upper management and the customer were now happy enough and now X.T. Products could go on to the next stage and see "if they could actually make this valve."

Design Engineering Work: Getting the Prototype to Pass "Spec"

As head design engineer for the OVRV, Vaughn had to make sure that the current design—in this case a working

prototype—functioned according to the “spec” on the print. This story differs from the others because Vaughn asked me to help him get this prototype up and running.

Vaughn got to the building room quite early to test and then attach an OVRV to a 1995 sport-utility gas tank that was to be shipped to the Arizona desert and run through the paces. A failure in the desert could seriously jeopardize a \$20 million deal with the customer and additionally compromise a step towards gaining the Vapor Recovery business at the other two domestic automobile manufacturers.

As is always the case in any product development setting there are constant time pressures. For Vaughn, he was supposed to ship the tank out today, but had only received the tank last night for attaching and testing. It was one of those high-pressure situations that has to be handled with patience, a willingness to try almost anything, and begs a little luck. The problem that he was having was that once he mounted the valve in the gas tank, the float (a small piece of plastic that rides within a slot in the body of the valve to insure a complete seal so gasoline will not escape, if the vehicle overturns) would not provide continuous seal at the proper level of pounds per square inch (or PSI). Vaughn quickly had to make a decision: He could either tear the valve apart to shave off something here or add a rubber gasket there or he could shave down part of the tank to create a more level attachment and, most importantly, produce successful test results. He looked at me and said, “I think I will try a rubber gasket first. Why don’t you hand me that sheet of rubber on the work table.”

Manufacturing Engineering Work: Building a High-Volume Machine

Even though Mick worked with some of the similar objects (i.e., drawings and parts) that the design engineers did, under the high-volume demands of his work, those parts take on a very different personality. This is a story about designing a high-volume assembly process and the challenges, problems, and worries that go along with it.

For Mick, one of the trickiest parts of building the OVRV is the variable orifice float subassembly. The challenge here is that the float has to provide a variable seal against the orifice that leads to an external vapor canister that holds and then condenses excess vapor from a gas tank. The challenge of the EPA requirement of a variable orifice sealing capacity is that the weight of the float and the strength of the spring have to be calibrated so that this variable sealing capacity is achieved at different PSI specifications. Of course, the problem is that not only does a shipment of 50,000 floats vary slightly in weight, but more frequently, a spring’s strength (gram load) can differ from spring to spring.

Mick had to design and build a machine that would assemble the float and the spring, place it in the body of the valve, recalibrate the spring with the weight of the float, and then test the subassembly 10,000 times a day, 3,000,000 times a year, or more to the point, one every 2.5 seconds. The critical challenge to this 2.5-a-second cycle time is recalibrating the strength of a

spring as it relates to the weight of the particular float. However, before this critical issue could be dealt with, an automated assembly process had to be developed. One of the perennial problems in automating a small mechanical valve is the likelihood of the spring being loaded improperly. Although a final pressurized test would detect 99% of those problems for Mick, he would forewarn me that having “large amounts of scrap lying on the floor is way too expensive for me (both in terms of cost and reputation).”

Production Work: Getting Product “Out the Door”

As a production technician, Jim’s job is to keep the assembly line moving. This is a story about that very thing—having to deal with the problems caused by all of the objects (i.e., parts and machines) that everyone else has designed, and with the added requirement of guaranteeing “on-time” delivery.

As a supplier in the automotive industry, on-time delivery is sacred. With the delays that have been occurring on the valve, Jim’s and his supervisor’s necks were on the line. Sure, his supervisor could blame the part’s bad design, a bad machine, or even Jim, but when it comes to on-time deliveries, as a production shift manager it is his problem to solve. His supervisor demanded and then pleaded, “Jim, I’ve got to have 25,000 shipped by the weekend. I’ve got 10,000 in shipping, so that’s 15,000 in just over two days.” That meant Jim needed to get to maximum production in four to six hours if they were going to make it.

Jim obviously felt the pressure, but his frustration was that he had been trying to understand why the small flappers were sticking together and jamming the equipment for a week now. He had (1) adjusted the feeder bowl vibration, (2) replaced the grasping portion of the robotic arm, (3) placed a pneumatic tube in the bowl to provide puffs of air to untangle the sticky parts, and (4) even sorted the flappers 100% to insure that the problem wasn’t part-related. His only thought was to try to pinpoint a nonmechanical problem or an environmental cause. Jim had already removed all of the dirt and dust from Stations 4 and 7. His next thought was to begin to worry about temperature. It was winter now, but the recently installed temperature control system maintained a constant temperature year round. He still didn’t have any solid ideas. What he did next was start again at square one. But for Jim that would take precious time, “and that’s time I don’t have,” he would say shaking his head.

These vignettes give a clear glimpse that *within* each function knowledge is hard won and “at stake,” and so it does have consequences in working *across* practices. The “right numbers” may be very difficult to design. Hitting “spec” requires adding parts that may be very difficult to assemble. Meeting a “cycle-time” requirement in the lab may be hard to maintain, let alone guarantee on the production line.

In most new product development settings, all required functions cannot fully localize, embed, and invest their

Table 1 Summary Table of Objects and Ends Across Practices

Sales Work	
Objects	Ends
<i>Numbers:</i> Price/cost, specifications, volume requirements, and delivery dates.	"The Right Numbers." Must be negotiated amongst multiple interests.
<i>Paper:</i> Numbers, contracts, spreadsheets, faxes, and "bluebooks."	"Close the deal." Get and keep the business.
<i>Technology:</i> Computer analyses, databases, and faxes.	"Numbers set in stone." Stand behind the numbers that work.
Design Engineering	
Objects	Ends
<i>Drawings:</i> Sketches, prints specifications, and tolerances.	"Design Review." Prints pass design check. "A functioning prototype." Trial and error to hit required specifications.
<i>Parts:</i> Materials, cost, prototypes, gaskets, epoxy, tools, etc.	"We have a prototype that passes spec." Stand behind the prototype that works.
<i>Technology:</i> CAD system and testing equipment.	
Manufacturing Engineering	
Objects	Ends
<i>Drawings/Prints:</i> Design prints, assembly prints, machining, or assembly process design prints.	"Keep it simple stupid." Fewest parts and fewest operations—less to go wrong.
<i>Parts/Product:</i> A variety of parts, raw stock, and finished stock, gauging and testing specifications.	"A high-volume process." A highly automated process that meets cycle time and volume requirements.
<i>Technology/Machines:</i> Machining and assembly equipment, machining operations, fixtures, parts testing, cycle time.	Improving the "manufacturing process." Incremental fine tuning and dealing with breakdowns.
Production Work	
Objects	Ends
<i>Parts/Products:</i> Millions of parts: springs, raw stock, finished stock, parts, and tested products.	Products "out the door." Meet production schedules and on-time delivery.
<i>Schedules and charts:</i> Production schedules, today's volume requirements, scrap, and defects rates.	"Scrap rate has to conform." Reduce waste and limit defects.
<i>Machines:</i> Machine utilization, breakdowns, and on-going trouble spots.	"Keep the people [process] working." Fix machines, adjust operations, and make changes to the "process."

knowledge *within* practice at the same time. For example, a design engineer will be able to deliver a working prototype that passes spec to the customer before a manufacturing engineer will be able to design and build a manufacturing process to assemble the product. Thus, not all sources of difference and dependence are known up front; consequences will continue to emerge as the product development process evolves. Using a funnel analogy as a way of thinking about the product development process, by the time a downstream function (i.e., manufacturing engineer, production manager) begins to understand how upstream decisions make their objects and ends more difficult to problem solve, it is hard for a downstream function to make changes to that upstream knowledge. Further, as individuals problem solve and create new solutions within their practice, novelty is created which generates new dependencies and often negative consequences that have to be worked out over time *across* practices. Given the product development funnel, the activities or interfaces to change upstream knowledge have to recognize the challenge is often one of jointly transforming knowledge at a pragmatic boundary.

The purpose of this discussion has been to specify the conditions at a pragmatic boundary that successful activities and technologies have to be able to resolve. When a manufacturing engineer sees the consequences that the current design has on creating an effective "high-volume process," they need to be able to transform the design engineer's knowledge localized, embedded, and invested in a working prototype before it is delivered to the customer. To understand what such a transformation process is like, let's see how Mick and Vaughn dealt with their differences about whether to change the design to accommodate subassemblies.

"Across"-Boundary Vignette: Mick, Vaughn, and the Question of Subassemblies

One of the challenges that Mick faced at previous design review meetings was that the assembly drawings he brought to the meeting were not up to date. X.T. Products had recently moved to a new CAD system and the draftsman trained in the new system was working on a tight deadline to deliver a set of design drawings so prototype parts could be ordered and sent to the customer. Because of this, there hadn't been time to update the assembly drawing.

The value of the assembly drawing for Mick was that it reflected issues that were of concern to a manufacturing engineer—orientation of parts, their order, and the location of "sticky" parts—so his arguments about potential

assembly, testing, and quality problems could be more easily represented to others. A design drawing, on the other hand, only has a two-dimensional quality, but represents well the critical tolerances, functional specifications, and overall dimensions of the design. When Mick presented his concerns about assembly and scrap rates, suggesting the move to subassemblies to deal with them, the design engineers had a hard time finding the assembly drawings useful because they did not reflect their concerns—the current design. For Vaughn and the other design engineers, what exactly would be changed and the consequences could not be made clear in their discussions with Mick. The designers did not dismiss his points outright, but in the end did not go along with Mick's subassembly proposal because it demanded changes that could certainly affect the functional flows that are so critical to the OVRV's current performance. Besides, subassemblies had never been needed before.

By the time a newly hired draftsman got up to speed on the new CAD tool and produced an assembly drawing that reflected the current design, there were only eight weeks left until the prototype parts deadline. In the next OVRV design and review meeting, Mick made the same points, but now with an up-to-date assembly drawing that reflected the current design. Mick's arguments were no different from the ones he had made before (i.e., difficulties in assembly and testing that would occur), but they had a very different impact now that the assembly drawing reflected the specs, tolerances, and locations of critical sealing surfaces that were "at stake" for the design engineers. This allowed Mick to describe in detail the type of assembly process that would be required for the current design, specifically the drawbacks of building a process to assemble and test such a complex product at high volume without producing tremendous amounts of scrap.

It was at this point that Mick roughed in the groups of parts that he would put into each subassembly on another copy of the assembly drawing. The changes Mick made on the print identified four testable subassemblies, where they would be attached, and the approximate areas where design changes would be required to accommodate four "large" subassemblies. After this example was penciled in on the assembly drawing, vigorous discussions began about the benefits and drawbacks of changing the design. One of the drawbacks uncovered was the current way of attaching parts (a high-velocity spin weld that melts plastic surfaces together) had always resulted in warping parts the size of these subassemblies. Vaughn was now open to using four subassemblies—but an alternative way of attaching them had to be found quickly.

Two days later at the next meeting, Mick and Jerry (production supervisor and team member) had a proposal.

In a small-volume product produced last year, a "snap-fit" design had been used to connect large plastic surfaces with great success. They explained to the team that the "snap fit" had two advantages over the spin weld. First, the snap fit did not involve either a rapid spin weld of plastic parts or the warping and alignment problems that were common in large parts. Second, most of the snap-fit holes and their clips could be placed on the outside of the pump and should not degrade the current functionality of the valve in a significant way. Over the following week, two meetings, and several informal conversations around the CAD system, the OVRV was redesigned with four independently testable subassemblies with "snap-fit" holes and clips to provide the connection.

Analyzing the "Across"-Boundary Vignette

Mick had proposed going to subassemblies in meetings before; what was different this time? In this case, the objects, or more specifically the assembly drawings used by the group in the design and review meeting, were different. The assembly drawing that was not current did allow Mick to begin to represent his concerns about assembly and testing, but because it did not reflect the current design and concerns of the design engineer, it failed to be a useful tool in specifying the differences and dependencies between them. Without being able to specify his differences clearly, Mick could not begin to specify the consequences of how the current design would result in significant costs when the design went into production. With the up-to-date assembly drawing, however, both Mick and the design engineers were able to represent what was "at stake." For the design engineer, it represented the exact location of critical sealing surfaces, the order and "stack-up" of dimensions, and the current design as embodied in a working prototype. For the manufacturing engineer, all parts are represented three dimensionally, including the order and identification of "sticky" parts that create the challenge of building and testing a product in the cycle time allotted.

Having represented and specified their differences, the assembly drawing also represented their dependencies and their consequences (e.g., how the current design makes "scrap rates" high for Mick or how going to subassemblies might undermine the OVRV's current functional capability). This supported a process where the group could define a shared problem with the OVRV and begin transforming their knowledge (the current design) and accommodating new knowledge (four subassemblies with snap-fit holes and clips). In terms of objects and

ends, we can see that the updated assembly drawing provided a shareable object and the creation of shareable ends at the boundary, where before the objects used actually reinforced the boundary.

Boundary Objects in New Product Development

The difference between the drawings used in the design review was that in the first two cases the drawings functioned as “within”-practice objects, but the updated assembly drawing functioned as an “across”-practice or “boundary object” (Star 1989, Carlile 1997). The concept of a boundary object, developed by Star, describes objects that are shared and shareable across different problem solving, contexts. In Star’s (1989) study of heterogeneous problem solving, she observed that in spite of the tremendous differences between scientists in various disciplines, they nevertheless were often very successful in cooperating to create “good science.” She describes boundary objects as objects that work to establish a shared context that “sits in the middle” (Star 1989, p. 47). Below I adapt Star’s four categories of boundary objects (repositories, forms, and labels, ideal type or platonic object, terrain with coincident boundaries) to describe the objects and their use by individuals in the settings that I observed.

First, *repositories* (i.e., cost databases, CAD/CAM databases, parts libraries) supply a common reference point of data, measures, or labels across functions that provide shared definitions and values for solving problems. Repositories function advantageously as a shared resource from which to compare across different functional settings when doing cross-boundary problem solving.

Second, *standardized forms and methods* provide a shared format for solving problems across different functional settings. Forms come in a mutually understood structure and language (i.e., standards for reporting findings, problem-solving methods [8-D Forms, D-FMEA, P-FMEA], engineering change forms, etc.) that makes defining and categorizing differences and potential consequences more shareable and less problematic across different settings.

Third, *objects or models* are simple or complex representations that can be observed and then used across different functional settings. Objects or models (i.e., sketches, assembly drawings, parts, prototype assemblies, mock-ups, and computer simulations) depict or demonstrate current or the possible “form, fit, and function” of the differences and dependencies identified at the boundary.

Fourth, *maps of boundaries* represent the dependencies

and boundaries that exist between different groups or functions at a more systemic level. Maps (i.e., Gantt charts, process maps, workflow matrices, and computer simulations) help clarify the dependencies between different cross-functional problem-solving efforts that share resources, deliverables, and deadlines. Because of the similarity of Categories 3 and 4, I combine them into one category I call objects, models, and maps to streamline the discussion that follows.

Characteristics of “Effective” Boundary Objects

We learned from the example of Mick and Vaughn that not every object used works as a boundary object. A critical question that has not been addressed in the literature is: What is the difference between a good and a bad boundary object? Even more paradoxically, a method or object that worked as a boundary object in one setting can become a boundary roadblock when taken to another setting. Based on a collection of over 65 observations of using different boundary objects (i.e., drawings, prototypes, D-FMEA, process maps, etc.) in cross-functional settings, I identified three characteristics of a tool, method, or object that made them useful in joint problem solving at a given boundary.

First, a boundary object *establishes a shared syntax or language for individuals to represent their knowledge*. In the case of Mick and Vaughn, the shareable quality of the assembly drawing as a representation is enhanced because both parties are familiar with it. For Vaughn, the assembly drawing represented critical tolerances and functional specifications. For Mick, the assembly drawing provided a more three-dimensional representation of the orientation of parts and critical issues for assembly and testing. The design drawing was not an effective boundary object because what is “at stake” for both the design and manufacturing engineers could not be represented on it.

This first characteristic of a boundary object has much in common with the insights from the syntactical approach to a boundary. In that perspective, the importance of having a shared language or syntax to deal with a boundary is fundamental. In the case of a boundary object, a shared syntax or language of representing knowledge at the boundary is a required characteristic for dealing with any type of knowledge boundary. Table 2 provides a comparison across type of boundary, category of boundary object, and the characteristics of a boundary object to represent the relationships across these key concepts. Some shared syntax is an essential feature of all three categories of boundary objects. In the case of repositories, if the meanings of the words used to search and store knowledge are not shared, then what is retrieved

will hurt more than it will help. Standardized forms use a shared language of representing problems (i.e., D-FMEA). As in the syntactical approach, when novelty arises, a semantic boundary is faced and another characteristic is needed.

An effective boundary object at a semantic boundary provides a concrete means for individuals to specify and learn about their differences and dependencies across a given boundary. A concrete method allows individuals to specify what they know—what they worry about—as concretely as possible to the problem at hand. In the case of Vaughn, the up-to-date assembly drawing allowed him to specify his concerns about important specs and critical sealing surfaces. For Mick, it allowed him to specify the challenges of assembling and testing a complex product at high volume. To deal with a complex knowledge boundary, the differences and dependencies between functions or groups must be specified, and the up-to-date assembly drawing identified the critical dependencies between the “functional spec” approach of the design engineers and Mick’s concerns about assembly and testing.

“Standardized forms and methods” and “objects, models, and maps” have this second characteristic (see Table 2). For example, a D-FMEA (Design Failure Mode Effects Analysis) provides a structured space where the representatives from each function can specify their specialized concerns about the current design. The methodology behind the D-FMEA requires a cross-functional team to assign values and critical priorities to the consequences identified by the individuals from each function. Using physical prototypes in cross-functional problem solving highlights the literal value of a concrete object in specifying “functional” relationships amongst parts, but also the dependencies among parts that impact assembly and testing issues. Further, the tangibility of physical parts allows for an ease in specifying differences and dependencies; their value becomes clear as they anchor the “scenarios” told by individuals about possible trade-offs to pursue. The nature of the problem determines what is adequate concreteness for a given boundary object. For example, a “process model” is certainly less concrete than a physical part, but when it is used to represent and learn about the sources of a design “bottleneck” in a complex product development process, its particular “concrete” means suit the nature of the problem faced. Of course, once this specifying and learning of differences and dependencies has taken place, we are often left with negative consequences that must be resolved.

At a pragmatic boundary an effective boundary object facilitates a process where individuals can jointly transform their knowledge. If there are negative consequences identified, then the individuals involved must be able to

alter, negotiate, or change the object or representation used (i.e., Mick’s knowledge about assembly and testing problems before the updated assembly drawing). If an individual cannot transform the current approach to a cross-functional problem, their knowledge will have very limited impact in a product’s development. In the case of Mick and Vaughn, dealing with the consequences identified required them to propose alternatives (i.e., subassemblies) and then alter the knowledge used to define the design (i.e., location of four subassemblies) and the particular manufacturing process developed (i.e., snap-fit holes). Individuals must be able to draw on, alter, or manipulate the content of a boundary object to apply what they know and transform the current knowledge used at the boundary.

“Objects, models, and maps” are the only category of boundary object that directly supports transforming knowledge (see Table 2). Not only are these types of boundary objects the most helpful in dealing with pragmatic boundaries, they are also the most complicated and expensive to establish. However, what should also be recognized is that all three categories of boundary objects have a portfolio effect; repositories and standardized forms support the use of objects, models, and maps as well as support processes to manage knowledge at a pragmatic knowledge boundary. Further, the knowledge transformed and created through the use of objects, models, and maps can then be used to enhance the content of shared repositories and the use of standardized forms and methods.

What also should be recognized is that boundary objects are no “magic bullet” because their characteristics are hard to sustain as problems and people change. For example, a CAD model can be an effective boundary object at one stage, but can falter when taken to another setting where a key functional group cannot represent their knowledge or alter the current knowledge with a CAD model. Mick summed up the challenging characteristics of a boundary object to me without knowing it when he commented after a difficult meeting that “CAD can be an effective communication tool in one meeting, then a ‘bludgeoning tool’ in the next.”

The Role Boundary Objects Play in New Product Development

The role that boundary objects play in new product development is that they help establish a “boundary infrastructure” (Bowker and Star 1999) or “boundary process” (Carlile 2002) that individuals use to manage knowledge across a given boundary. In the case of the updated assembly drawing, it aided Mick, Vaughn, and others in representing their knowledge, learning about their differences and dependencies, then jointly transforming current

Table 2 Type of Knowledge Boundary, Category, and Characteristics of Boundary Objects

Types of Knowledge Boundary	Categories of Boundary Objects	Characteristics of Boundary Objects
Syntactic	Repositories	Representing
Semantic	Standardized Forms and Methods	Representing and Learning
Pragmatic	Objects, Models, and Maps	Representing, Learning, and Transforming

and more novel knowledge to resolve the negative consequences identified at the boundary. Many authors (Star 1989, Henderson 1991, Wenger 1998) have recognized the importance of boundary objects; other authors in the product development literature have recognized the value of prototypes and modeling in driving innovation (Leonard-Barton 1995, Clark and Wheelwright 1995, Iansiti 2000, Schrage 1999). This research connects these two different literatures and adds value to both by specifying different categories of boundary objects in new product development and the critical characteristics that are essential in establishing effective boundary processes.

What we see in this examination of the capacity of a boundary object is two-fold: both practical and political. Practical because it must establish a shared syntax or a shared means for representing and specifying differences and dependencies at the boundary. Political because it must facilitate a process of transforming current knowledge (knowledge that is localized, embedded, and invested in practice) so that new knowledge can be created to resolve the negative consequences identified. This practical and political capacity of a boundary object at a pragmatic boundary provides an *infrastructure* or *process* where current and more novel forms of knowledge can be jointly transformed, producing more shared knowledge or syntax at the boundary.

Conclusion

This research demonstrates at a deeper level why communication across functional boundaries (Dougherty 1992) is hard, given the problematic nature of knowledge in practice. Further, it also describes and proposes what can be done to support individuals who need to work across such boundaries. This allows us to examine more closely what lies inside such conceptualizations as “adequate” information processing capacity (Galbraith 1973). What we see at a pragmatic knowledge boundary

is not just a matter of processing more knowledge, but processes for transforming knowledge. This research has also provided an opportunity to revisit how knowledge and boundaries are conceptualized in organizations—and by extension a fundamental aspect of organization theory—differentiation and integration (Lawrence and Lorsch 1967). Knowledge and differentiation were reframed as knowledge in practice and differences in kind. The challenge of integration (Grant 1996) was concretely described in the relation between the type of boundary faced and the category and characteristics of the boundary object used. Below I summarize these relations and link them to the notion of “integrating devices” (Lawrence and Lorsch 1969) to clarify the contribution of this research and the work that still needs to be done.

A *syntactical* approach is based on the existence of a shared and sufficient syntax at a given boundary. A sufficient syntax is efficient because differences and dependencies have been specified and agreed to in advance. From this approach *integrating devices* are inherently syntactical “processing” tools (i.e., shared repositories, taxonomies)—and integration is accomplished through processing information or transferring knowledge across a boundary. However, when novelty arises the sufficiency of the syntax is in question and another boundary is faced.

A *semantic* approach recognizes that differences exist or emerge overtime, so individuals have different interpretations of a word or an event. In this way, the semantic approach recognizes that there are always differences in kind (i.e., thought worlds, Dougherty 1992) and the emergence of novelty on one or both sides of the boundary is a natural outcome in settings where innovation is required. From this approach *integrating devices* should be seen as processes or methods (i.e., standardized forms and other shared methods) for translating and learning about the differences and dependencies at a boundary. In some cases a process of translation is sufficient, but when negative consequences are identified another boundary is faced.

A *pragmatic* approach recognizes that differences in knowledge are not always adequately specified as differences in degree or interpretation, but that knowledge is localized, embedded, and invested in practice. This pragmatic framing of knowledge highlights the negative consequences that can arise given the differences and dependencies at a boundary. To resolve these consequences this paper outlines a process of transforming knowledge—where individuals represent, learn, negotiate, and alter the current knowledge and create new knowledge to resolve the consequences identified. Here, *integrating devices* recognize that knowledge has to be transformed; to create new knowledge, old knowledge has to be changed (Teece

et al. 1997). We see this with the example of the up-to-date assembly drawing where knowledge was integrated as current and more novel forms of knowledge where jointly transformed at the boundary.

This pragmatic view of knowledge and boundaries has been helpful in explaining why knowledge is both a barrier to and a source of innovation in a product development setting. This view has also been used to clarify the role that boundary objects play in establishing an infrastructure or process where knowledge can be represented, learned, and transformed. Further, as organizations become more specialized and society grows more complex, this research is a reminder that the challenge is not just one of processing capacity, but also representational capacity at a boundary. The value of taking a pragmatic view of knowledge and boundaries is that it moves us closer to recognizing and focusing our research on the challenges of knowledge representation in organizational life.

Acknowledgments

The author would like to thank the Special Editors, Arie Lewin, Anna Grandori, Bruce Kogut, and the two anonymous reviewers for their dedicated help. He would also like to thank the following colleagues who provided suggestions that vastly improved this paper: Lotte Bailyn, Rebecca Henderson, Starling Hunter, Rakesh Khurana, Carsten Osterlund, Eric Rebentisch, John Van Maanen, and Eleanor Westney. This work was supported in part by the MIT Center for Innovation in Product Development under NSF Cooperative Agreement Number EEC-9529140.

References

- Allen, T. 1971. Communications, technology transfer, and the role of technical gatekeeper. *R&D Management* **1** 14–21.
- . 1977. *Managing the Flow of Technology*. MIT Press, Cambridge, MA.
- Ancona, D., D. Caldwell. 1992a. Demography and design: Predictors of new product team performance. *Organ. Sci.* **3** 321–341.
- , ———. 1992b. Bridging the boundary: External activity and performance in organizational teams. *Admin. Sci. Quart.* **37** 634–665.
- Ashby, R. 1956. *Introduction to Cybernetics*. Wiley, New York.
- Bertalanffy, L. 1956. General system theory. L. Bertalanffy and A. Rapoport, eds. *General Systems: Yearbook of the Society for the Advancement of General Systems Theory*, vol. 1. Society for General Systems Research, Ann Arbor, MI, 1–10.
- Bourdieu, P. 1977. *Outline of a Theory of Practice*. Cambridge University Press, Cambridge, UK.
- , L. Wacquant. 1992. *An Invitation to Reflexive Sociology*. University of Chicago Press, Chicago, IL.
- Bowker, G., S. Star. 1999. *Sorting Things Out: Classification and Its Consequences*. MIT Press, Cambridge, MA.
- Brown, J. S., P. Duguid. 1991. Organizational learning and communities-of-practice. *Organ. Sci.* **2**(1) 40–57.
- Brown, S., K. Eisenhardt. 1995. Product development: Past research, present findings, and future directions. *Acad. Management Rev.* **20**(2) 343–378.
- Buckley, W. 1968. Modern systems research for the behavioral scientist. W. Buckley, ed. *Modern Systems Research for the Behavioral Scientist*. Aldine Publishing, Chicago, IL.
- Carlile, P. 1997. Transforming knowledge in product development: Making knowledge manifest through boundary objects. Unpublished dissertation, University of Michigan, Ann Arbor, MI.
- . 2002. An integrative framework for managing knowledge across boundaries. Working paper, MIT Sloan School of Management, Cambridge, MA.
- Clark, K., S. Wheelwright. 1995. *The Product Development Challenge: Competing Through Speed, Quality, and Creativity*. Harvard Business Press, Cambridge, MA.
- Cook, D. N., J. S. Brown. 1999. Bridging epistemologies: The generative dance between organizational knowledge and organizational knowing. *Organ. Sci.* **10**(4) 381–400.
- Dougherty, D. 1992. Interpretive barriers to successful product innovation in large firms. *Organ. Sci.* **3**(2) 179–202.
- Fleck, L. 1935/1979. *Genesis and the Development of a Scientific Fact*. T. J. Trewn, R. K. Merton, eds. and F. Bradley; T. J. Trewn, trans. University of Chicago Press, Chicago, IL.
- Galbraith, J. 1973. *Designing Complex Organizations*. Addison-Wesley, Reading, MA.
- Grandori, A. 1986. *Perspective on Organization Theory*. Ballinger Press, Cambridge, MA.
- Grant, R. 1996. Prospering in dynamically competitive environments: Organization capability as knowledge integration. *Organ. Sci.* **7** 375–387.
- Harper, D. 1987. *The Nature of Work: Working Knowledge*. University of Chicago, Chicago, IL.
- Henderson, K. 1991. Flexible sketches and inflexible data-bases: Visual communication, conscription devices and boundary objects in design engineering. *Sci., Tech., & Human Values* **16**(4) 448–473.
- Iansiti, M. 2000. *Technology Integration: Making Critical Choices in a Dynamic World*. Harvard Business School Press, Cambridge, MA.
- James, W. 1907. *Pragmatism*. The American Library, New York.
- Jantsch, E. 1980. *The Self-Organizing Universe*. Pergamon Press, Oxford, U.K.
- Joyce, W. 1986. Matrix organizations: A social experiment. *Acad. Management J.* **3** 536–561.
- Katz, R., M. Tushman. 1981. Project performance and the locus of influence in the R&D matrix. *Acad. Management J.* **28** 67–87.
- Keller, R. 1986. Predictors of the performance of project groups in R&D organizations. *Acad. Management J.* **29** 715–726.
- Kogut, B., U. Zander. 1992. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organ. Sci.* **3**(3) 383–397.
- Lave, J. 1988. *Cognition in Practice: Mind, Mathematics, and Culture in Everyday Life*. Cambridge University Press, Cambridge, U.K.
- , S. 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.
- , ———. 1969. Organization-environment interface. *Developing*

- Organizations: Diagnosis and Action*. Addison-Wesley Publishing, Reading, MA, 23–30.
- Leonard-Barton, Dorothy. 1995. *Well Springs of Knowledge: Building and Sustaining the Sources of Innovation*. Harvard Business School Press, Cambridge, MA.
- Nonaka, I. 1994. A dynamic theory of organizational knowledge creation. *Organ. Sci.* **5**(1) 14–37.
- , I. Takeuchi. 1995. *The Knowledge-Creating Organization*. Oxford Press, Oxford, U.K.
- Orr, J. 1996. *Talking about Machines: An Ethnography of a Modern Job*. Cornell University Press, Ithaca, NY.
- Peirce, C. S. 1898/1992. *Reasoning and the Logic of Things*. Harvard University Press, Cambridge, MA.
- Polanyi, M. 1966. *The Tacit Dimension*. Anchor Day Books, New York.
- Redding, C. W. 1972. *Communication Within the Organization*. Industrial Communication Council, New York.
- Reddy, M. J. 1979. The conduit metaphor. A. Ortony ed., *Metaphor and Thought*. Cambridge University Press, Cambridge, U.K.
- Reynolds, B. 1987. *Material Anthropology: Contemporary Approaches to Material Culture*. University of California Press, Berkeley, CA.
- Schrage, M. 1999. *Serious Play: How the World's Best Companies Simulate to Innovate*. Harvard Business Press, Cambridge, MA.
- Shannon, C., W. Weaver. 1949. *The Mathematical Theory of Communications*. University of Illinois Press, Urbana, IL.
- Spradley, J. 1979. *Ethnographic Interview*. Holt, Rinehart, and Winston, New York.
- . 1980. *Participant Observation*. Holt, Rinehart, and Winston, New York.
- Star, S. L. 1989. The structure of ill-structured solutions: Boundary objects and heterogeneous distributed problem solving. M. Huhns and L. Gasser, eds. *Readings in Distributed Artificial Intelligence*. Morgan Kaufman, Menlo Park, CA.
- Suchman, L. 1987. *Plans and Situated Actions: The Problem of Human Machine Communication*. Cambridge University Press, Cambridge, U.K.
- Szulanski, G. 1996. Exploring internal stickiness: Impediments to the transfer of best practice with the firm. *Strategic Management J.* **17** 27–43.
- Taylor, C. 1992. To follow a rule. . . . C. Calhoun, E. LiPuma, and M. Postone, eds. *Bourdieu: Critical Perspectives*. University of Chicago Press, Chicago, IL.
- Teece, D. J., G. Pisano, A. Shuen. 1997. Dynamic capability and strategic management. *Strategic Management J.* **18**(7) 509–533.
- Tushman, M. 1978. Information processing as an integrating concept in organizational design. *Acad. Management Rev.* **3**(3) 613–624.
- von Hippel, E. 1988. *The Sources of Innovation*. Oxford Press, London, U.K.
- , M. Tyre. 1996. The mechanics of learning by doing: Problem discovery during process machine use. *Tech. and Culture* **37** 312–329.
- von Krogh, G., K. Ichijo, I. Nonaka. 2000. *Enabling Knowledge Creation: How to Unlock the Mystery of Tacit Knowledge and the Release the Power of Innovation*. Oxford University Press, New York.
- Wenger, E. 1998. *Communities of Practice: Learning, Meaning and Identity*. Cambridge University Press, Cambridge, U.K.