

ORGANIZATIONAL CAPACITIES FOR SUSTAINED PRODUCT INNOVATION

Deborah Dougherty
Associate Professor
Faculty of Management, McGill University
1001 Sherbrooke St. W.
Montreal Que H3A 1G5
email: doughert@management.mcgill.ca

I: Definitions and Conceptual Background

This report is a very preliminary summary of the results of a study on the organization-level capacities that enable an organization to generate multiple new products or services over time. The intent is to develop the basis to assess the innovativeness of an organization, understand what managers need to change to improve their innovativeness, and identify how they can manage that change. This paper emphasizes the theory that underlies the phenomenon of “organizational capacities for innovation.” Before developing the theory on capacities for innovation, however, it is essential to define terms, since at least some of the conflicts and incongruities in the innovation management literature arise because people are simply talking about different things.

In this paper, “product innovation” refers to the whole process of bringing a new product or service to the market, and includes product conceptualization, design and development, production or operationalization, marketing, distribution, and selling. Some think of “innovation” as the creative idea generation or technological invention parts of the process. However, these are limited views, since any or all of the steps in the innovation process may require creativity, and the most innovative aspect of a product may be the novel integration of apparently standard steps.

“New” refers to uses, users, technology, marketing, distribution or logistics, and any combination of these *that are unfamiliar* to the organization. The first use of new-to-the-world inventions is a very limited subset of product innovation, so focusing only on them does not include most product innovation in most organizations. “New to the firm” as a definition of innovations includes imitations. However, a strategy of sustained innovation based on imitation is rarely effective (Freeman 1980; Mitchell and Singh 1992) so the definition does not over-emphasize imitations. Product innovations also range from incremental to radical, but research

suggests that many established organizations cannot develop even incremental innovations on a sustained basis, so excluding incremental innovations from our attention may also be very limiting (Fast 1978; Henderson and Clark 1990; Dougherty and Corse 1995).

“Sustained product innovation” involves the simultaneous development of multiple product innovations, while also managing mature businesses efficiently. Single innovation projects have been researched extensively, and how to manage single projects effectively is well understood (Rothwell 1977; Cooper 1983; 1988; Kanter 1988; Souder 1987; Johne and Snelson 1988; Clark and Fujimoto 1991; see reviews by Brown and Eisenhardt 1995; Dougherty 1996). Most organizations can produce a successful new product or service occasionally, so an emphasis on how one-off innovations can be accomplished is not particularly new. As firms must adapt on a more continuous basis to changes in markets, technologies, and competition, however, occasional new products may not be enough to remain viable. *Sustained* product innovation poses major organizational challenges that are not addressed by a focus on single project management. Sustained innovation means that the organization must: (1) provide resources to, monitor, and staff a variety of different projects at different stages of development all at once (including routine and mature products), (2) develop procedures so that all projects both use well-established steps for effective innovation project management and adjust these steps to their unique situation, (3) enable all these projects to build on the firm’s core competencies, (4) enable the core competencies to build on the new knowledge and insight created through specific innovations, and (5) manage the firm’s resources such as manufacturing, R&D, selling, and distribution systems so that they both accommodate the new products readily and run as efficiently as possible.

The ability to do all of these things is the problem that this paper addresses. Given the importance of this problem, it should come as no surprise that hundreds of answers and part answers exist. The innovative organization is a problem solving, market oriented, self-organizing, learning, n-form, transnational, front-back organization that thrives on chaos (Clark and Fujimoto 1991; Day 1990; Weick 1977; Hedberg 1981; Bartlett and Ghoshal 1991; Galbraith 1995; Peters 1988). All of these insights are most likely right, somehow, but not all of them connect directly to the messy complexities of sustained product innovation, or to one another. Any effort to clarify the organizational capacities for sustained product innovation should build

on these insights but also overcome two critical problems: (1) how can non-innovative organizations change; and (2) what do we mean by “organization” anyway?

The first problem is the failure to deal with change. Many organizations do not use the myriad tools, techniques, perspectives, and approaches for being innovative, despite the widespread dissemination of these ideas (Cooper and Kleinschmidt 1986; Majahan and Wind 1992; Dougherty and Heller 1994). The years-long struggle among firms in many industries (autos, communications, airlines, electronics, utilities) to adapt themselves to more innovation indicates that becoming and staying innovative is no easy thing. Moreover, studies indicate that the innovative organization is not less organized and more chaotic as some believe, but rather it is based on a fundamentally different system of organizing (Burns and Stalker 1961; Jelinek and Schoonhoven 1990; Dougherty and Corse 1996). To become capable of sustained product innovation may require changing the very approach to managing and organizing, which is a far more significant adjustment than adding new processes or new performance measurements. To get beyond this problem, the development of yet another description of the innovative organization will not do. It is important to analyze capacities for innovation in such a way that we can also understand how to change from non-innovative to innovative

The second problem is a lack of consensus concerning what “organization” and “capacity” are in the first place. This problem arises no doubt because the fields of innovation and technology management cut across many disciplines, from physics to psychology, from economics to sociology, from operations research to accounting. The diversity of views is less like blind men touching different parts of an elephant, however, and more like people at an accident scene who literally see something different. I propose that a sociological view is a useful initial basis for analyzing organizational capacities for sustained product innovation, because it directly addresses both the complex actions and relationships that are inherent in innovation, and the problems of *becoming* and then *being* capable of sustained product innovation. This is, of course, just one perspective, but it address the social and human aspects of organizing that are an important part of sustained innovation.

A sociological view defines the organization as a complex social system (or set of social systems), the daily functioning of which is affected by socially shared norms, routines, and roles and relationships, and by status groups, and interest groups, and procedures for authority and

control. These shared social understandings and structures can either enable or impede the myriad activities of innovation such as making sense of customers, understanding the goals, forming teams, solving technical problems, and so forth. Organizations can also develop strong inertial forces from a combination of existing routines, premises, and role expectations, and studies demonstrate consistently that it is necessary to change the *whole system* to effect major organizational change (not necessarily everybody's understanding of everything all at once, but sets of understandings that cut across the organization – see Bartunek 1983; Barley 1986; Westley 1990; Reger et. al., 1994).

Unfortunately, two views of organization which ignore its complex, social system nature dominate the field of innovation management. One is the “best practices” approach, which concentrates on the various tools, strategic views, techniques, processes, and procedures (Wheelwright and Clark 1992; Griffin and Hauser 1993; Cooper 1988). These tools and techniques are most likely correct in describing what innovative organizations *should* do, but they do not explain how they, as social systems, can do them. A related emphasis holds up “hi-tech” firms as the exemplars for all innovative organizations, even though most organizations are not “hi-tech” in the sense that they operate in an exploding, rapidly developing technology regime. In fact, organizing non “hi-tech” firms to be innovative may be more difficult, since they must combine the management of complex technologies with complex marketing techniques and mature products. The question must be, what are the underlying organizing capacities that enable people to use these techniques and to manage rapidly developing technologies effectively? Presuming that change follows readily from outlining the most rational or optimal course of action belies much of the empirical record regarding the management of change.

A second limited view of organization is what I call the “plumbing and wiring” view. This view focuses on organizational structure in the abstract, and simplifies complex organizational processes into such things as lines and levels, degree of formalization and decentralization, and number, length, and width of communication channels. This very mechanical view of the organization presumes that the complexities of everyday working and thinking can be readily manipulated by adding more “channels” or eliminating levels. Research shows that these conceptions of organization have weak, no, or unexpected associations with effective product innovation (see Downs and Mohr 1976; Nord and Tucker 1987; Damelpour

1991), suggesting that they are, at best, gross simplifications of the social behavior they represent. One case study of five firms in electronics finds that they are in fact more organized, with more units, levels, and procedures, than non-innovative firms, not less (Jelinek and Schoonhoven 1990). The real difference is that people in the innovative firms changed those structures all the time, because they had another understanding of how to work that allowed them to manipulate these structures readily to fit changes in work.

An interpretive view of the organization deals more directly with the shared understandings people use to frame work. This view focuses on how people collectively make sense of their day-to-day world using “interpretive schemas” (Bartunek 1984), which, like culture, are shared understandings and assumptions about why events happen as they do and how people are to act in different situations. They arise through the collective solution of organizational problems like who are we, how do we work together, how do we deal with authority (Schein 1990). Interpretive views have been used in a variety of ways to understand the complexities of innovation, and so fit the problem well. For example, Dougherty (1990; 1992a) explains problems forming multi-functional teams by suggesting that departments are like different thought worlds: they both know different things about customers and products, and know things differently, so problems with innovation are not simple matters of “communication flow” or goal conflict. Nelson and Winter (1982) argue that organizational routines provide the occasion for innovation by highlighting deviations from routine. Barley (1986) demonstrates how the same new technology can alter organizations differently, depending on how new roles are negotiated.

Conceptualizing the organization as a “negotiated order” provides a useful image of organization for the purposes of understanding capacities for sustained innovation (Strauss 1971; Brown 1992; Orlikowski 1994). “Negotiated order” emphasizes the way in which the patterns of behavior, including structures of roles and relationships, are continually recreated by the actions and interactions of their members. For example, as people attend meetings, work at their lab bench, propose an advertising approach, and otherwise operate in the daily organizational context, their job standards, roles, and relationships with bosses and co-workers are in effect “negotiated.” This is not to say that the social order is created anew everyday, since the existing social order conditions subsequent social negotiations. Thus, particular patterns of managerial

styles, communication, and emphases develop and remain because they “work” and become taken for granted. The organization’s negotiated order is conditioned further by the need to generate more revenues than costs, the firm’s technologies and markets, and governmental regulations for employment, accounting, etc. Variation arises across organizations and across time even under the same constraints, because rules are incompletely understood, used as a resource, stretched and broken, and reproduced with only indifferent fidelity.

The outcome is that during periods of stability people pretty much know what to do every day when they come to work, what to expect from others, how things are evaluated, and generally how to “get on,” because the enormous complexities of social action “settle down” into common sense and shared understandings. However, when conditions change significantly (e.g., the “sudden” appearance of new competition or technology), these negotiated orders no longer work, and the myriad complexities of everyday work are again up for grabs. People become anxious when pressed to do things for which they have no commonly agreed-upon rules or which violate their understanding of the organization’s identity, which is one reason why change is so difficult. In the first major study of organizing for innovation, Burns and Stalker (1961) found that many of the electronics firms in their sample did not become innovative because people did not know how to organize for it – they kept re-instituting variations on bureaucracy even though these did not work, because that was the only way to organize that they knew.

To apply the organization as negotiated order to innovation, we first have to break the negotiated order down into pieces that can be separated conceptually (the strong popularity of the “best practices” and “plumbing and wiring” views of organization is due in no little part to the fact that they both translate readily into simple survey questions). Patterns that are negotiated around innovation can be thought of as “cultural capacities” for organizing (Swidler 1986; Dougherty and Corse 1995). A cultural capacity is a “tool kit” of cultural material such as symbols, stories, habits, categories, informal know-how and skills from which people draw to develop a line of conduct in a certain realm of action. We all have cultural capacities for such things as parenting and family life, being a citizen in a democracy, and being an employee. If one stops to think about any of these roles, it should be clear that they are complexes of thinking and acting that are socially negotiated and reinforced. Organizational capacities for sustained

innovation are like conceptual maps of the terrain; there are different kinds of maps for different purposes, and different groups of people become accustomed to using certain maps over others.

The next step is to identify sets of activities that need to be carried out for innovation, and for which capacities will have emerged to map them. Looking at how these activities are in fact carried out in a range of organizations, from non-innovative to innovative, indicates different capacities for the same set of activities, some that enable innovative practice, and some that disable it. I propose three sets of activity: knowledge management, task arrangement, and the management of human relations.

In the next section I first summarize the general best practices in each set of activities, outline the organizational capacity found in my research that allow innovative organizations to carry out these complex activities effectively, and contrast that with the capacity operating in non-innovative organizations. A more detailed discussion of all three capacities for organizations with different degrees of innovative ability, along with a fuller examination of how organizations change from one kind of capacity to another will be developed later, as the research on which this summary is based proceeds.

Section II: Innovation Activities and the Supporting Organizational Capacities

The organizational capacities described here have been developed from a study of about 11 different organizations, and the ideas may be limited by the kinds of organizations included in the research. The study methodology and sample are described in the appendix.

Rather than finding a clear distinction between innovative and non-innovative organizations, the results suggest that there is a continuum of innovative ability, from non-innovative to comprehensively innovative. Non-innovative organizations create new products outside the normal organizational system, in special task forces or venture units. The organizational capacities of the normal system are geared to support the continued functioning of the system, and so enhance the production of established products and services rather than new ones. Near the middle of the continuum, organizations create new products easily, but only in their familiar product categories. Here, the organizational capacities are geared to support a focus on the product category and on “doing business” in an MBA-like fashion, with senior managers emphasizing optimization of production and “portfolio management” in the abstract, and

business and product line managers emphasizing everyday product development processes. Toward the more innovative end of the continuum, organizations generate a variety of product innovations, from incremental to radical, and manage them rather easily in conjunction with existing businesses. Their organizational capacities focus people's attention on creating value by solving customers' problems rather than on a specific product category or internal functioning in the abstract. Senior managers develop competencies and manage the organization as a social system, middle ones manage the businesses to generate profits, and lower level ones create products.

Three activities and their associated capacities are described separately, but the reader will notice that all three act together.

A. Creating, Organizing, and Using Knowledge for Innovation

Product innovation depends heavily on developing and applying knowledge in three distinct domains: (1) about markets and customers, (2) about product and process technology, and (3) about the organization's own core operating abilities. Aspects of all three kinds of knowledge must be connected in the specific product's design and business plan, in the guiding strategy for the business, and in the organization's policies and procedures governing technology, manufacturing, and other major systems in its value chain. How well organizations can develop each domain fully and integrate all of them would have a lot to do with how well they can innovate.

Customer and Market Knowledge: First, the largest single cause of new product failure is the failure to understand customer needs (Rothwell et al 1974; Cooper 1983; 1988). To put it more positively, successful new products have more market knowledge than failed ones (Souder 1987; Mohrman 1996). However, customer needs for new products are difficult to articulate and sort out, because the customers themselves may not be sure, the knowledge itself is often tacit, and that knowledge may be embedded in a context of use and not retrievable except by hands-on interaction with that context (von Hippel 1994). Customer needs may change in any case as the product is used (Rosenberg 1982). Despite these difficulties, the more details on various needs and attributes from customers the better for a product's ultimate success (Bacon et. al, 1995). This should make sense, since any product is in fact a complex bundle of components, attributes, and specifications, and the more the many design and engineering choices that go into its construction are based on actual customer priorities and preferences, the better (see Griffin and

Hauser 1993 on QFD, a technique based on the generation of at least 200 specific need statements, and Leonard-Barton 1991 on empathetic design, or extensive interaction and prototyping with users).

These deep, “visceral” insights on the customer and the context of use must be complemented with knowledge from other levels of analysis (Dougherty 1992b). Assessments of the size of the market, likely segments, how quickly is the market is likely to evolve, and what are the key underlying drivers of that evolution are essential, as is knowledge of more general trends in demand, government regulations, competition, and complementary products that might affect the size and growth rates of the potential market (Gatignon and Robertson 1985). As well, organization members must understand the products’ connection to existing product lines and brands, since a bad fit can stretch abilities or damage the reputation of the firm, but too close a fit can keep the organization from appreciating new opportunities (Christensen and Bower 1993; Day 1990). None of these levels and aspects of customer/market knowledge can be ascertained for sure, but it should be obvious that designing a product with undesirable performance attributes for a small set of potential users who probably won’t change current patterns of use anyway is *not* a good thing.

Technology Knowledge for Product Innovation: Customer knowledge must be complemented with technology knowledge before a product can be created (Freeman 1982; Cohen and Levinthal 1990). “Technology” is used to refer to science or engineering-based knowledge used to create products and manufacturing processes, such as microbiology, materials science, information systems, software, chemical engineering, or electronics. Like the first domain, knowledge of technology ranges in level from the specific details of a particular technical problem to systems of technologies that may underlie a given product line. Since technologies are often part of a larger system, the ability to solve any particular problem may depend on access to the broader knowledge of the field or discipline, regarding what progress is being made where in solving various puzzles (Freeman 1982; Rosenberg 1982). To access such knowledge, organizations are advised to develop their own technologies systematically, and to create “platforms” of capability so that some of the problems that come up in an evolving product line have already been worked on (Sanderson and Uzumeri 1996; Meyer and Lehnerd 1997). They are also advised to monitor emerging technologies that may ultimately replace theirs (Tushman and Anderson 1986; Utterback

1992). However, the ability to grasp these apparently systematic trajectories and interconnections, especially in the earlier phases of a technology's emergence, requires hands-on experience, learning by doing, and extensive informal social networks (Nelson and Winter 1976; von Hippel 1988).

Core Operations Knowledge for Product Innovation: Knowledge of core operations refers to the knowledge of what the organization as a system can do, of how various sets of technologies can be combined with other organizational systems to create and launch products – or, knowledge of how to make things happen. This domain of knowledge is reflected in Leonard-Barton's (1992) managerial systems, in Quinn's (1994) set of intellectual or service skills that are central to the intelligent enterprise, and Henderson and Cockburn's (1995) architectural competence, which is the ability to integrate technology components into effective product packages. Knowledge of core operations has a particular, hands-on aspect as described by Brown and Duguid's (1991) discussion of "communities of practice." A community of practice is a set of people who draw on each other's ideas, past experiences, and collective interactions to make the tricky interpolations and strategic improvisations that are necessary to solve the complex problems of implementing sophisticated machines and processes. In communities of practice, problems are situated in the actual activity, and people "...See the way in which the process of doing the task is actually structured by the constantly changing conditions of work." Knowledge of core operations also operates at more general levels, so people understand the business well enough to appreciate how to add value (Jelinek and Schoonhoven 1990), and it accumulates into systematic "learning curves" (Adler and Clark 1991).

All three domains of knowledge are vital to the success of any particular product, new or old. In content, all three run the gamut from minute, situated details to general structures, trends, and trajectories that operate outside the organization. And all three are strongly based on experience, on having deep, intimate understanding of the issues in their contexts of customer use, technological field, and organizational potential.

Innovative Versus Non-Innovative Capacities for Knowing:

My findings suggest that innovative organizations have a capacity for organizing knowledge – what I call a "knowledge architecture" – that allows them to both know more about each domain of knowledge, and to integrate all that knowledge more readily in the management of all activities,

including product innovation. Descriptively, the knowledge architecture allows the organization to differentiate each of the three domains from each other, thus highlighting each one and elaborating the knowledge within each domain by level. The differentiation and elaboration are very important, because they allow the organization as a social system to not only recognize the separate importance of each domain, but also to actively develop each as a competence, apart from specific businesses or product lines. Capital and research investments are made for long term development of the competence, not only for short term business opportunities. The knowledge architecture also contains numerous integrative mechanisms at all levels of the organization, so that knowledge of markets, technologies, and core operations can be easily combined and re-combined into new products, new businesses, and new competencies. New insights in any domain can enter at any level and be woven into the overall knowledge. People are skilled at “mixing and matching” technical and operational abilities to create a product, and the whole system is geared toward continual “creative destruction” (Shumpeter 1950), in which resources are reorganized and re-configured to meet new needs and opportunities.

The capacity to know more is not simply a matter of creating more categories for storage and use, and more coordinating devices. The innovative knowledge architecture uses three qualitatively different frames of reference or socially shared and understood framing devices that people can draw on to invoke, make sense of, and apply various aspects of the three knowledge domains. These frames of reference are like lenses (or, in “net” parlance, “search engines”) that people can collectively develop and use to create a product, manage a business, and re-create a strategy. The contextual nature of these frames of reference is fundamental. Each frame deals directly with actual contexts of practice: the context of the organization’s relationship with its industry and market, the context of the problem-solving abilities of each domain of knowledge, and the context of the customers’ problems that products must solve. I borrow an analogy on imitation from Warner (1991) to explain the importance of real contexts as opposed to abstractions of them. He points out that a man pretending to be a woman can know only so much about women because he has never lived life as one. His knowledge is “grainy,” while a woman’s knowledge of herself is smooth and dense. Similarly for innovation, non-contextual frames of reference provide only grainy knowledge; for example only marketing people or top managers have real-world knowledge

of customers, and technology people may have extensive knowledge of metallurgy theory but limited appreciation for how it can fit into certain real-world problems for users.

The three frames of reference focus on what should be known, how it should be known, and how the knowledge should be used. The organization's shared understanding of its relationship with its industry and market, called the organizational identity, frames what should be known. An organization's identity refers to what people perceive as central, enduring, and unique or distinctive about their organization and believe others share as well (Weick and Westley 1996, citing Dutton and Penner 1993; Albert and Whetten 1985). It is an image that aids sensemaking, according to Ring and Van de Ven (1994:100): "By projecting itself onto its environment, the organization develops a self-referential appreciation of its own identity, which in turn, permits the organization to act in relation to its environment." The actual identity, like culture, is unique to an organization, but the kind of identity in innovative organizations highlights all three domains of knowledge. The organization's members see the firm as a problem solver for customers, as a community of practice that applies its knowledge to work out customer problems. People define what they are good at collectively as making customers successful with their objectives, which provides a vivid, active image of them in the context of working with customers, and which justifies and makes sense of all three domains of knowledge. This kind of identity enables people to see beyond product categories or functional boundaries to relationships with customers. Perhaps counter-intuitively, this frame of reference focuses attention on the actual value creating opportunity and on what the organization can do.

The second contextual frame of reference connects the three separate knowledge domains by enabling people to break down the complexities of each with the same conceptual categories that are connected to real-world problems. The domains are differentiated into basic building blocks of product creation which are both sensible to the domain and connected to customer expectations. All the domains are anchored with external referents that make sense across domains. This approach to framing the domains allows people to both appreciate the kinds of knowledge development that should be made within the domain, and to relate the separate parts of knowledge to one-another across domains. Domain knowledge is thus substantive, or focused on how it contributes to the problem solving activities of the firm as a whole. For example, all domains have the same categories and goals, such as reducing costs for customers, improving customer

satisfaction in particular ways, meeting commonly understood performance goals, and particular knowledge is generated around solving those problems.

The third frame of reference anchors people's attention on the customer's actual problem situation. People are skilled at seeing the situated realities of customers' operations or lives; they can step through mediating frames of "product category" to not only see but make sense of the complex world which is the customers' contexts.

The key to being able to have and use more knowledge for innovation, therefore, seems related to people's skills and experiences with using these context-based frames of reference to invoke knowledge that cannot always be articulated or transmitted. Rather than have everybody know everything, everybody has some experience in aspects of two of the three frames. Individuals work usually from both a substantive domain and in a customer or market context, so they all have some substantive knowledge of how the products are made and applied knowledge of how products are used. When they work together on an innovation project or organization-wide task force they can readily negotiate a shared frame of reference for the work, and a common language for the task. Then each person applies his or her own expertise to the commonly appreciated task. People from different disciplines can both understand the expectations put on the product, and how their own expertise can contribute to those expectations. They do not have to understand each other's knowledge. People are thus able to create many more products and see how these products fit together in a business, because they do not have to spend the time translating expectations into marching orders for each department, and then re-translating output into common abstractions so senior managers can appreciate the implications.

In contrast, the dominant logic of the non-innovative organization's knowledge architecture is simplification, or the reduction of complexity. All knowledge is broken down into separate categories that concentrate on just one area, and filter out extraneous "noise." The identity of the non-innovative organization is an abstract state of being (e.g., we are a university, or a pillar of the community), and the only common focus people have is an internal one: on the ongoing functioning of the organization as an operating system. People understand what they are good at in terms of their own function. Operating knowledge in the abstract dominates, and the other domains are seen as subordinate subsets of operations. Domain knowledge is broken down further in a way that makes sense only within the domain, as part of the logic of simplification by

separation and abstraction. So, various parts of different knowledge domains can only be put back together by experts in that domain, for the purposes of the domain's role on the ongoing functioning of the system. Knowledge is developed within the domain to solve problems of ongoing functioning, as those problems have been translated into standards and objectives for that function or area. The emphasis on abstraction inhibits people's ability to see the situated realities of the many customer contexts, and instead focuses them on how the product should be used rather than how it is used or why.

The non-innovative organization has a very complex knowledge architecture, with highly developed capacities for generalizing and separating, and concentrating on developing one's expertise for keeping the overall system running. People are not only trained in most higher or professional education programs to think this way, they are encouraged to develop these capacities further by most jobs in many large organizations. But the very logic of this architecture of knowledge inhibits the collective ability to break things down and reconfigure them, which is what innovation is all about. The lack of differentiation among knowledge domains means that the competencies of understanding customers and markets, and of developing technologies to meet, or indeed create, opportunities are always under-developed. Market knowledge is nothing more than marketing knowledge, or techniques to launch and sell regular output. Technology knowledge is also just another part of operations, a step in the design of products or the improvement of cost ratios for the system. Long term investments in any are not made, because they do not make sense – people cannot see the outcome. Insights about new markets or technologies just do not make much collective sense unless they relate to the usual products. Only senior managers can “see” the whole, because the knowledge from the different domains integrates back up only at the top. Senior managers therefore tend to dictate the products or product categories while others execute their specialized knowledge.

Only a few organizations in my study operated with such a knowledge structure in full force because of its inability to enable new products. They have added marketing departments and have learned to use rigorous market research (rather than senior management hunches) in evaluating opportunities and designing new products. A number are still fixated on ongoing functioning of the system, however, so that keeping the network, plants, or logistics systems running “at optimum” is the primary focus. The focus continues to highlight internal operations and functioning over

external changes and shifts. Without the collective ability to frame things in contextual terms, people cannot focus their knowledge on value creation.

The real break with the non-innovative knowledge architecture comes when another domain of knowledge is differentiated and understood to be the equal of operating knowledge, not just a part of it. Empirically, the best example occurred in a consumer products company that “suddenly” lost market share because competitors were offering new products that better met the changing market needs. All three frames of reference shifted to some extent to incorporate the contexts. First, it shifted its identity from an internally defined one of “our food is best because it is,” to one defined externally by the product category and brand: “we make the best food in this segment, for these people.” Their customer and product frame changed from how the product should be used (it is the best food so of course everyone will love it) to how it actually is used (this is the best food for this use). The domains of marketing and product technology were strengthened around application of their knowledge in context, but the manufacturing group still concentrated on non-contextual, self-referential goals (e.g., line speeds, similar packaging). This results from the lack of differentiation of technology as a substantive domain of knowledge, or as a competency for the long term. New investments are tied to specific new business plans, and are not be made unless the proper pay-back from the business can be demonstrated – which it often can’t since specific product ideas are too small to finance major technology changes. Therefore, only incremental improvements are made in technology, so only incremental product adjustments are made.

The transition to the more comprehensively innovative knowledge architecture depends on the full differentiation of all three knowledge domains, and their development as competencies that can be drawn on for continual innovations. Both the innovative and non innovative capacities for organizing knowledge do not stand alone, however. They require profoundly different capacities for organizing work, and for defining and maintaining the relationships between the individual and the organization.

B. Organizing the Work of Innovation

All of the knowledge of innovation must be applied to the work of innovation. Knowing implies doing so these activity sets are two sides of the same coin, but elaborating on the basic tasks of sustained product innovation assures that the capacities cover all necessary social action.

Product innovation, in general, is a process. The process includes various tasks such as opportunity identification, concept development and verification, product design, manufacturing, marketing, and distribution, and ongoing management. These tasks usually cannot be separated or even run sequentially for innovation, since they tend to flow together, and often proceed in parallel, so that the carrying out of one task must take into account the carrying out of another (Souder 1987; Cooper 1994).

For example, Clark and Fujimoto (1991) show how Japanese auto producers reduce time to market by working the product design and manufacturing design in parallel rather than in sequence. Tooling is ordered and production processes set out while the product is being designed, which can only happen when the two groups work very closely, with in-depth appreciation for one another's problems and constraints, and extensive, rich communication. Marketing may be investing considerable resources in developing the new product's image in the marketplace based on certain features, which requires additional mutual adjustment across departments if those features change. Therefore, people working on a specific innovation must be able to recognize and solve complex problems quickly. They must push issues along within their own area using extensive expertise, but also jointly focus on problems that affect more than one function or unit and solve those problems by taking each other's constraints into account. Yang and Dougherty (1993) call this approach to work "iterative organizing." The same kind of iterative organizing must go on throughout the organization, among products and business units, to assure that all products draw on competencies and resources and fit in the competitive strategy, that the organization can leverage investments in technologies that are too large for any one product or business, and that the administrative processes and procedures themselves are continually revised to accommodate new activities (Jelinek and Schoonhoven 1990).

Organizing the work of innovation is not just a matter of putting in teams, adjusting the number of levels, or even "thriving on chaos" (Peters 1988). Studies indicate that, rather than being non-organized and chaotic, innovative organizations have lots of structures and procedures, so that people always know what they are supposed to do and to whom they can go for help or approval. The difference is that these structures and procedures are adjusted easily as needed to fit the specific work (not for vague goals of cost-cutting or "effectiveness"). Structures and processes are understood as tools rather than devices for marking out turf or controlling behavior.

The question becomes: how do people conceptually frame and make sense of all these activities so that their organization systems “organize” but change as needed?

Innovative Versus Non-Innovative Capacities for Organizing Tasks:

My study suggests that innovative organizations have a fundamentally different collective task architecture which provides a very different mapping for who does what work and how than that in the non-innovative organization. This task architecture allows people to work both laterally and vertically in the iterative organizing mode, solving problems, but still maintain the integrity of the process of product innovation. The innovative organization has the capacity to organize its work as a hierarchy of processes that preserves the whole of the product innovation process, focuses people’s attention on problem identification and solving. The non-innovative organization has the capacity to organize its work as a hierarchy of separate tasks, which preserves the integrity of each task in the abstract so that each can be carried out in an optimal way, provided of course that the ongoing functioning of the system remains stable.

The innovative task architecture is characterized by three key themes: (1) the delineation of processes of work rather than steps in the process; (2) hierarchy based on different realms of responsibilities and authority rather than declining degrees of the same responsibility and authority, and (3) a focus on problem solving rather than on the application of established solutions.

First, the complex whole of work is broken down into smaller pieces so that people can comprehend their role, but the delineation of work is by processes rather than by separate tasks or steps in the process. The corporate, business unit, and product-line levels is each responsible for the entire process of value creation, from opportunity identification to sales, which preserves the integrity of the process. At each level, people can go back and forth from opportunity identification to technology development, for example, and adjust either or both as appropriate. The innovative organization is indeed a “process organization” as others recommend (e.g., Galbraith 1995; process re-engineering views, see Hammer 1995). However, the underlying capacity is people’s ability to understand their work as the whole process and to see their role in that whole. They can concentrate on the flow of work and on the inter-connections and iterations among steps, functions, and procedures, and rethink and re-negotiate specific steps and

connections as adjustments in the whole are developed. This attention to the process is reinforced by a wide variety of procedures for decision-making, operations review, process evaluations, assessments of customer satisfaction, product development processes, planning, and so on, which are regularly developed to articulate the processes. These procedures are evaluated and changed regularly, however, to make sure that they enable the product development process.

Preserving the integrity of the process gives everyone a similar understanding of the opportunity because they helped define it, and therefore a common view of the objectives of work. Because of the extensive delineation of processes, people know how the tasks connect, and are skilled at refashioning connections as necessary. People are specialists not just in a single area expertise, but in how their expertise connects to the product development process, however that process shifts over time.

Second, the processes are clumped into different levels, so each level in the organization, from corporate to business unit to product line, is responsible for a different set of problems in the overall process of delivering value to customers. The hierarchy is one of different responsibilities for different kinds of problems, and people at each level operate on the basis that people at the other levels are doing their work. For senior levels, the problems concern maintaining and developing the organization's knowledge competencies, so they make corporate-wide resource investment decisions for R&D, manufacturing technology, and selling and distribution systems that are shared by the businesses, they evaluate strategic opportunities of the firm in its industry or industries and make new long-term investments in competencies, and closely manage the organizational processes and procedures to assure that connections are working, and that the standards that are used to guide work fit the work. Senior managers develop, implement, and replace standards for work to keep the processes flowing. Business unit managers work on the problems related to keeping the business unit in sync with its market so that it can generate profits. Their job is to know the business well enough to forecast their needs in production capacity and selling and distribution well enough so that their colleagues managing the functions can allocate resources, and to make longer term predictions of where the business is going so that senior managers can invest in the competencies. They also must reconfigure the business unit to accommodate new products, and organize and manage systems that integrate the

various tasks in product development. Product line people work on creating new products, managing the product line, and developing new opportunities for it.

The third aspect of the innovative task structure arises in part because of the delineation of processes over separate tasks, and of responsibilities by type of problem rather than by office. It is a strong focus on solving problems. Problems regularly arise in the execution of the whole processes of innovation, and people in this task structure are skilled at identifying and working on them. The actual solutions are not pre-ordained, so the problem solving is very creative.

The innovative organization is indeed highly organized and extensively controlled. All the processes developed at each level articulate aspects of the overall process of value creation, and each level works in the context of the others. Decisions are based on enhancing processes rather than functions in the abstract. So the development of new manufacturing capability or IS resources are done as part of the processes of business management or product line development. Product innovation is just part of the normal work of the organization, no different than other work, and in fact helps the organization to uncover glitches in connections.

These are very different capacities for organizing tasks than those in the non-innovative organization. Again, this capacity in the non-innovative organization is equally complicated, but based on a fundamentally different logic, that of “scientific management.” First, the different tasks of product development are themselves differentiated both hierarchically and laterally. The task of opportunity identification (or definition) is done at the top of the organization by senior managers. They determine the situation for the organization as a whole, and convert that determination into standards which are expected solutions of the grand problem. These become output objectives or marching orders for each of the functions or business divisions, used by them to frame their task or piece of a task. The tasks in the whole process are separated so that each can be developed, examined, and adjusted to be as efficient or otherwise optimal as possible without having to worry about ramifications on other parts. The logic presumes that the whole remains stable, and that each part both can and should be optimized separately. The connections between the parts are minimized and simplified, each unit or function is a repository of solutions to partial problems, and people are highly skilled at applying these established solutions to new problems.

Only a few organizations in the study have the non-innovative task architecture in its complete form, since it cannot handle activities that are inherently coordinating, process-based, and lateral, like product innovation. However, a number of the organizations are not too far removed from it. Several have a business structure, but the “business unit” is really only the marketing function which carries out marketing tasks. The whole process is not managed deliberately as in the innovative organization, and business managers must negotiate access to technology, production, or other functional resources on an almost ad hoc basis, because the functions may still be managed separately to be “optimal” (usually defined as limited cost or less cost than last year). In one firm, for example, the sales function called for a moratorium on new products because the salespeople had to be “pulled of the streets” for several days for each introduction, and they were losing too many selling days a year. In the innovative organization, the whole selling process is geared to regular introduction of new products, the salespeople are familiar with the changing customer context, and the new products are designed with the selling needs in mind in the first place, so this problem is continually addressed.

The key for transition to the middle of the innovative ability continuum is to shift everyone’s attention onto the business market, and to highlight the importance of working with customers directly. Doing so is facilitated if the unit is organized by whole business unit which contains all the separate tasks and functions necessary for product development, and if it is organized within by product line process, not function. The strict business unit structure is a temporary development step, I think, perhaps necessary because it cuts out a “bite-sized” whole with its own subset of development processes that people can comprehend. The product line organization within allows managers to collapse functions into one another, reducing the barriers. The organizations in my study at least retained their functional definitions of task above the product line level, which separated the whole back into the tasks, and encouraged managers to focus on optimizing their tasks. As well, the corporate managers still focused on abstract tactical standards like returns on invested capital (defined in a particular way), volume, scrap rates, and other throughput productivity measures. The transition to the more comprehensively innovative capacity for organizing tasks begins to occur when senior managers take on the responsibility for managing the organization – making sure that the connections work, that resources flow properly, that decisions are made quickly, and that the standards used to guide work fit the work.

C. Managing the Work Relations:

The final set of activities concerns the connections between individuals and the work organization. The work relationship centers on an explicit or implicit “contract” between the worker and the owner (now, usually “management”) regarding the kind and amount of work to be done in return for pay and other benefits associated with a livelihood, such as a career and social status. This relationship is in fact very complex, because even for the simplest of jobs no one can specify in advance exactly what is to be done, what is the role of the person and the rights and responsibilities of the employee and the employer. Instead, these can be worked out only in the continuing day-to-day relations. Work relationships also involve differential authority, super and sub-ordination, which heightens the conflicts of interest and concerns over rights and responsibilities. Central concerns in managing the work relations are how to define the tasks that each person is responsible for, how to coordinate different tasks, and how to assure that work meets standards, all in a way that is acceptable to both employees and employers.

The work of innovation can exacerbate problems with these concerns, because much of the day-to-day work cannot be explicitly specified, and the outcomes cannot always be measured cleanly. Employers must invest large sums of money without immediate verification or assurances that the resources will work, which may make them extra anxious to monitor people’s work. Employees must take on roles that are broad enough to enable them to work across the whole innovation process, and flexible enough to allow them to work with a variety of other people directly. They must have the skills to quickly form complex but often temporary working relationships based on mutual adjustment, and must feel responsible for making the innovation happen no matter what their part in it is. In other words, it seems that employees must do anything and everything at all times, and employers must make sure resources are always available without feedback on returns.

In their early study of innovation, Burns and Stalker (1961) found that innovative work was “experienced by the individual manager as an uneasy, embarrassed, or chronically anxious quest for knowledge about what he should be doing, or what is expected of him, and similar apprehensiveness about what others are doing” (1996 ed., p. 122). Innovation was inhibited in the firms they studied because many people refused to work that way, and concentrated instead

on their own interests and career. Without an organization capacity to frame and make sense of this work relationship so that both sides are comfortable and understand their roles and responsibilities, people will fixate on specifying their rights and responsibilities.

Innovative Versus Non-Innovative Capacities for Work Relations:

The innovative organizations in my research defined the work relationship in a fundamentally different way, one that supported and enhanced innovative work rather than inhibited it. People were able to quickly negotiate simple, comfortable connections with one another and work closely together. Three themes characterized the capacity for innovative work relationships: (1) the personal nature of the relationships, (2) the expectation that people would work to make the task interesting for one another, and (3) the understanding that the primary impetus for action was the “law of the situation,” that everyone was authorized to grasp the opportunities.

First, people relate to each other on work projects on a personal basis, so that the relationships are first hand rather than distant, multidimensional rather than simple, and rich rather than abstracted as is the case with limited role play. Personal relationships are capable of more mutual adaptation, in which the work emerges as the situation demands through the reciprocal adjustments of the people involved, but they are also more intense and require careful mutual attention to the feelings and integrity of the other person. It is not that employees devote themselves heart and soul to the work organization, in some sort of “total institution,” however. The other two innovative capacities allow people to handle these more intense work relationships because they significantly reduce the number of issues that people must negotiate to carry out a joint task. The practiced skills people have with working with customers, invoking common frames of reference, and applying their expertise to solving problems keep people focused on common issues, make individuals’ roles clear and simple to understand, and provide a clarity of purpose.

People do not have to negotiate specific roles, role expectations, task goals, and decision criteria among all the players every time collaboration is called for, nor worry that some people may change their outcome objectives later on. In addition, the organization’s many resources,

systems, processes and procedures are managed to support the product development process. The individual innovation teams do not have to figure out how to redesign the computer system, work out how engineering change orders are managed, or fix the evaluation procedure to enhance their own ability to work, since these systemic issues are being worked on by management.

The second theme is a particular norm of reciprocity: that people must strive to make an activity interesting for others to encourage them to join in. This means that the activity must meet some goals of the other person, and that the initiator takes some responsibility in seeing that that happens. As well, one cannot make work interesting for another unless one can really appreciate what the other wants or needs, and is able to listen effectively. The individual is free to choose to participate or not in any given project, within the boundaries of the job which says that everyone works on processes with others.

Finally, the central impetus for action is the law of the situation. The particular situation dictates what needs to be done for product and business development, and individuals are expected to respond to that situation to solve specific problems, take advantage of opportunities, make judgements, and balance the contradictions that may occur between new problems and established practices or standards. Again, the other two capacities focus people's attention on customers and business opportunities, and provide them with the shared skills to apply their knowledge and to work through problems. The law of the situation both recognizes people's responsibility for making things happen within certain realms of action and certain business areas, and gives people the authority to control aspects of the work situation. The organization as a whole frames the kinds of actions and opportunities that should be addressed and how, but the individuals involved have the authority to work out a specific action themselves.

In contrast, the capacity for work relationships in the non-innovative organization develops and reinforces secondary interpersonal relationships that are based on formal job descriptions, and informal expectations that build up around the occupation or function. Each individual is expected to carry out her or his own job properly in accordance with the standards set by the managers and is held personally accountable for doing so. Norms of reciprocity usually concern staying out of one another's turf. People do not develop skills with creatively negotiating work with others but rather with doing what they are told or (if in a professional or managerial position) what is expected of them. As people go higher in the hierarchy, they have

more and more autonomy to execute their office as they personally see fit, and everyone understands this as the boss's prerogative. Integrative work may conflict with their job expectations if it pressures them to compromise functional standards, or to make choices that are not seen as part of their jobs. There is no reason to make things interesting for others, because everyone does what he or she is supposed to do. The central impetus for action is to meet standards and especially to avoid mistakes that may bring the whole system down. Individual action and the individual herself are controlled and constrained, and ideally prevented from having conflicts or contradictions.

The non-innovative capacity for work relations fits well with the related knowledge architecture and task architecture, and reinforces the inability to innovate in a sustained way. One can see how people with such an understanding of work relations would find innovation work very anxiety-provoking. The organizations in my study which are still not fully innovative have developed the innovative capacities for work to varying degrees within teams at lower levels, simply because effective innovative team work requires the more personal and reciprocal kind of relationships. However, these relationships are not reinforced by the organization, and so are tenuous.

The real break from being non-innovative to being incrementally innovative is the formal institution of teams for managing new product development, coupled with granting the teams authority over all the design choices and product strategy. If senior managers retain the authority to dictate product design and if product approval continues to be a lengthy political process rather than a matter of strategy, then team members have no reason to invest themselves fully in the task. Indeed, their role is still one of execution or carrying out orders, and that is more easily done (at least in the immediate term) when people maintain functional autonomy. However, any functional separation in the task architecture over different steps maintains a sense of dominance which induces an opposing sense of separateness. Continuing to monitor manufacturing's compliance with abstracted, non-strategic through-put performance measures also reinforces non-innovative capacities for work relationships in these departments. In team meetings they will focus constantly on their constraints or what they cannot do, which inhibits creativity and the degree of innovativeness.

The transition from the incrementally innovative to the comprehensively innovative capacities for work relations depends, of course, on developing the other two capacities as well. People cannot engage routinely in the more intense personal relationships unless managers have begun to manage the organization as an innovative supporting system, unless working in processes is fully supported, and unless competencies are developed so that resources are available. I think a key is to stop defining work and work identities based on function. The areas of expertise are important and must be developed, but function should not define knowledge, task, and work relationships.

The data so far suggest that only some effort toward each of these changes is necessary to initiate the transition. For example, in one incrementally innovative firm, product teams have slowly taken on more authority simply to get their work done properly. Now they are beginning to attack more fundamental infrastructure problems in manufacturing and how those resources can be made more flexible – not because they have chosen to invade that turf, but because these problems must be solved if they are to get their new products out the door on time. At the same time, the corporation is beginning to integrate businesses which pulls together resources. What in fact may happen is that they may be ready, finally, to differentiate the technology knowledge domain and invest in long term technical capability, because managers of these larger business units will be able to “see” the payback. Corporate research (now run in separate units for tax purposes – which of course inhibits all three sets of capacities) has begun to put together world-wide task forces for bits of technology that many units can share (e.g., pasta making equipment). These resources should help managers develop their unit technologies, and ultimately enhance the creation of technology as a domain of knowledge, a competence, corporation-wide, rather than just a part of operations. These separate initiatives, started one by one because they attack various problems in maintaining the corporation’s competitiveness, can impel the organization toward becoming more capable of sustained product innovation. Provided, of course, that senior managers keep being innovative as a strategic goal, and take the lead in changing how they approach their work.

More work needs to be done on specifying and clarifying these capacities, and in describing in more detail how they can be developed in organizations with a history of stable, non-innovative operations. In particular it is necessary to expand on how managers actually

change, and how they combine new ways of acting with new ways of thinking. I am now working with a team of three doctoral students to analyze the data fully, to collect more data from very innovative organizations, and to verify these emerging ideas. However, I think this emerging theory on capacities for innovation is useful for three reasons. First, consistent with change theory, it acknowledges that whole sets of thinking and acting need to be changed at once if the organization as a social system is to become more innovative, and indicates what those systems are. It is therefore realistic. Second, it connects these whole systems directly to innovative (and non-innovative) activities, explaining both why some organizations can innovate and why some cannot. Third, while these have not been well developed yet in this preliminary report, the theory points to manageable dynamics for change. This approach should lead to a more complete, yet when developed, more elegant, understanding of the capacities necessary for an organization to be truly innovative, and how organizations can develop these capacities.

METHODS APPENDIX

I interviewed people working on new products, their functional and business unit managers, and corporate managers in 11 organizations. The organizations are established firms located in North America, come from a variety of mostly mature industries, and include service and manufacturing. This is a theoretical sample, and organizations were included provided they fit with the more mature nature of the industries yet added important variation around issues that the analysis suggested was important. The first criterion of inclusion was is the organization actively attempting to improve or develop its product innovation abilities, so the data are limited to those organizations for whom improved innovation was important, at least to some senior managers. A second criterion for inclusion was the organization's ability for sustained product innovation, since we sought a variety of abilities. However, I could not gain access to presumed exemplars of innovativeness such as 3M or Levi Strauss, so our sample is skewed to the non-innovative end of the continuum. I also deliberately added service firms with differing innovative abilities, to assure that our findings did not depend on service versus manufacturing. The actual unit of analysis is the business unit within the organization, although everyone interviewed also discussed the whole organization and its impact on innovation activities.

I used a common interview protocol in all organizations, with different questions for people at different levels. People who were working on specific innovation projects were asked to describe whether and how they incorporated customer, technology, and core operations knowledge in their design and development effort, how they get it, and what kinds of systems or processes existed in the organization to help them (or get in their way). Functional or unit managers were asked to describe how they make choices, allocate resources to various domains, evaluate projects, assist their project managers in acquiring resources and knowledge, and any changes that have occurred in the organization to promote innovation. Senior managers were asked to describe the strategies for innovation, how they developed resources and knowledge abilities, and any programs they used to improve innovation. People come from all functions and specialties.

Capacities represent shared, interpretive understandings for what knowledge, tasks, and work relationships are meaningful. The focus on the interviews was on people's rationales, how they interpreted others, and the meanings they assign to events, to elicit their interpretive understandings that reflect the knowledge architectures they operated with.

To measure the relative ability for sustained product innovation, I relied on people's descriptions of how readily and how well they felt that their business unit developed new products or services. Peoples' assessments were for the most part honest and fair, with some bias in both directions – some people dwelled on problems, while some glossed over problems that others we interviewed in the organization pointed out. Ideally I would use an objectively measured, industry comparison of new product success rates and proportion of revenues generated from new products introduced in the past 3 or 5 years, all by the degree of innovativeness of the product. Unfortunately such measures are not available. This less than ideal method for categorizing organizations allows us to compare the organizations

internally in these data, but also limits the inferences we can draw regarding the application of these findings to other organizations. For this analysis, I classified the 11 organizations into a low, medium and high sustained innovation categories. Table 1 summarizes the organizations, the number of people interviewed in each, and the category of sustained innovation. Five organizations are low in sustained innovation ability, three are near the middle, and three are high.

The data introduce two important limits on the kinds of inferences we can draw regarding knowledge architectures and how they affect product innovation. First, for the most part the organizations come from relatively more stable (historically) but non-commodity industries, where discontinuities have been less abrupt, changes less rapid, and markets less determined by global supply and demand than one might find in electronics, biotechnology, or, say, aluminum. The knowledge, task, and work relations architectures we find may relate only to organizations with similar experiences. Second, the data are cross-sectional, so we have no direct evidence of change. However, several of the organizations were in the midst of major change, while several others had made major changes in recent years, according to the people interviewed. I use these indirect data to infer some dynamics of transformation from one knowledge, task, and work relations architecture to another, but these inferences are extremely tenuous. Moreover, the developmental path that we infer no doubt relates to more established organizations, because it evolves from an internal fixation on operations, to the inclusion of customer knowledge, and ultimately to the anchoring of the knowledge architecture on core operating abilities. Organizations in newer, high technology industries may begin with a fully developed technology domain that may evolve with different dynamics.

Data Analysis: My research assistants and I compared and contrasted people's descriptions of how they made sense of and used the three domains of knowledge in product innovation, organized their work, and connected themselves and others, using methods described by Bailyn (1977) and Strauss (1987) to create a summary without imposing our own interpretations. We first analyzed several interviews from one company in depth, to highlight a variety of possible themes, in several different two hour sessions, with several different groups of research assistants. We continued these intensive open coding efforts by including more interviews in more firms, to elaborate and develop the issues until it seemed that no new issues were emerging (approximately 10 2-3 hour sessions, with two different groups of 3 or 4 people, to develop as much insight as possible). Next, we discussed key themes and identified those associated with the dimensions described below. We returned to the data in several sessions to discuss each of these key themes across all interviews, and then each author worked independently to summarize the themes across a subset of interviews. We met several more times to discuss the overall patterns of the three capacities that seemed to best explain the differences in the data.

We are now developing approximate indicators to capture the degree to which the key themes occurred in an organization. Students have already coded a subset of the interviews for knowledge, and other codes are being developed.

References (incomplete)

Adler, P. and K. Clark (1991) Behind the Learning Curves, Management Science, 37:267-281.

Bacon, G., S. Beckman, D. Mowery, and E. Wilson (1994) "Managing Product Definition in High-Technology Industries: A Pilot Study," California Management Review,

Bailyn, L. (1977) Research as a cognitive process: Implications for data analysis. Quality and Quantity, 11:97-117.

Barley, S. (1986) "Technology as an Occasion for Structuring," Administrative Science Quarterly, 31:78-109.

Burns, T. and G. M. Stalker (1961) The Management of Innovation. London:Tavistock.

Brown, J. and Duguid (1991) "Organizational learning and communities of practice," Organization Science, 2:40-57.

Brown, R. (1992) Understanding Industrial Organizations, London: Routledge

Brown, S. and K. Eisenhardt (1995) "Product Development: Past Research, Present Findings, and Future Directions," Academy of Management Review, 20:343-78

- Christensen, C. and J. Bower (1993) "Catching the Next Wave: Why Good Customers Make it Hard," working paper, Harvard Business School.
- Clark, K. and T. Fujimoto (1991) Product Development Performance, Boston: Harvard Business School Press.
- Cooper, R. (1983) A process model for industrial new product development. IEEE Transactions on Engineering Management, 30:2-11.
- Cooper, R. and E. Kleinschmidt (1987) Success factors in product innovation, Industrial Marketing Management, 16: 215-233.
- Cooper, R. (1994) "Third Generation New Product Processes," Journal of Product Innovation Management, 11:3-14
- Cooper, R. (1994) "Third Generation New Product Processes," Journal of Product Innovation Management, 11:3-14
- Day, G. (1990) Market Driven Strategy: Processes for Creating Value, New York: The Free Press.
- Dougherty, D. (1990) Understanding new markets for new products, Strategic Management Journal, 11:59-78.
- Dougherty, D. (1992) A practice-centered model of organizational renewal through product innovation, Strategic Management Journal, 13:77-92.
- Dougherty, D. (1992) "Interpretive Barriers to Successful Product Innovation in Large Firms," Organization Science, 3:179-202.
- Dougherty, D. and T. Heller (1994) The illegitimacy of successful new products in large firms. Organization Science, 5:200-218.
- Dougherty, D. and S. Corse (1995) When it Comes to Product Innovation, What is so Bad about Bureaucracy? Journal of High Technology Management
- Dougherty, D. and S. Corse (1996) What Does it Take to Take Advantage of Product Innovation? Research on Technological Innovation, Management and Policy, Rosenbloom, D. and R. Burgelman (eds), JAI Press.
- Fast, N. (1978) The Rise and Fall of Corporate New Venture Divisions, UMI Research Press.
- Freeman, C. (1982) The Economics of Industrial Innovation, Cambridge, MA: MIT Press.
- Galbraith, J. (1995) Designing Organizations;
- Gatignon, H. and T. Robertson (1985) "A Propositional Inventory for New Diffusion Research," Journal of Consumer Research, 11:849-867.
- Griffin, A. and J. Hauser (1993) "The Voice of the Customer," Management Science, 12:1-27.
- Hage, J. (1988) ed. Futures of Organizations, Lexington MA: Lexington Books.
- Henderson, R. and K. Clark (1990), Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms, Administrative Science Quarterly, 35: 9-31.
- Henderson, R. and I. Cockburn (1995) "Measuring Core Competence? Evidence from the Pharmaceutical Industry,"

- Jelinek, M. and C. Schoonhoven (1990) The Innovation Marathon: Lessons From High Technology Firms, Oxford: Basil Blackwell.
- Kanter, R. M. (1988) When a thousand flowers bloom, in B. Staw and L. Cummings, eds
- Leonard-Barton, D. (1991) Inanimate integrators: A Block of wood speaks, Design Management Journal, 2:61-67.
- Leonard-Barton, D. (1992) Core capabilities and core rigidities: A paradox in managing new product development. Strategic Management Journal, 13:111-126.
- Majahan, V. and J. Wind (1992) "New Product Models: Practice, Shortcomings and Desired Improvements," Journal of Product Innovation Management, 128-139.
- Meyer, M. and A. Lehnerd (1997) The Power of Product Platforms, New York: Free Press.
- Moorman, C. (1995) Organization market information processing: Cultural Antecedents and new product outcomes, Journal of Marketing Research,
- Nelson, R. and S. Winters (1982) An Evolutionary Theory of Economic Change, Boston: Belknap Press.
- Nonaka, I and I. Takeuchi (1995) Knowledge Creating Organizations, NY: Oxford University Press.
- Nord, W. and S. Tucker (1987) Implementing Routine and Radical Innovations, Lexington, MA: Lexington Books.
- Peters, T. (1992) Liberation Management: Necessary Disorganization for the Nano-Second Nineties; New York: Knopf
- Quinn, J.B (1994) The Organization as an Intelligent Enterprise, Academy of Management Executive.
- Reger, R., L. Gustafson, S. Demarie, and J. Mullane (1994) "Reframing the Organization: Why Implementing Total Quality is Easier Said Than Done," Academy of Management Journal, 19:565-584
- Rosenberg, N. (1982) Inside the Black Box: Technology and Economics, Cambridge: Cambridge University Press.
- Rothwell, R. (1977) "The Characteristics of Successful Innovators and Technically Progressive Firms," R&D Management, 7:191-206.
- Sanderson, S. and M. Uzumeri (1996) Innovation Imperative, Burr Ridge, IL: Irwin Publishing
- Souder, W. (1987) Managing New Product Innovations, Lexington, MA: Lexington Press.
- Strauss, A. (1987) Qualitative Analysis for Social Scientists, New York: Cambridge University Press.
- Tushman M. and P. Anderson (1986) "Technological Discontinuities and Organizational Environments," Administrative Science Quarterly, 31: 439-465.
- von Hippel, E. (1994) "'Sticky Information" and the Locus of Problem Solving: Implications for Innovation," Management Science, 40:429-439.
- Von Hippel, E. (1988)
- Warner, W. (1991) "What Computers Don't Know," book review, Technology Review, May-June 67-69.
- Wheelright, S. and K. Clark (1992) Revolutionizing Product Development, New York: The Free Press.

Yang, E, and D. Dougherty (1993) Product Innovation: More Than Just Making a New Product, Creativity and Innovation Management, 2:137-155.