



Organization Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Deliberate Learning and the Evolution of Dynamic Capabilities

Maurizio Zollo, Sidney G. Winter,

To cite this article:

Maurizio Zollo, Sidney G. Winter, (2002) Deliberate Learning and the Evolution of Dynamic Capabilities. Organization Science 13(3):339-351. <https://doi.org/10.1287/orsc.13.3.339.2780>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

© 2002 INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Deliberate Learning and the Evolution of Dynamic Capabilities

Maurizio Zollo • Sidney G. Winter

Department of Strategy and Management, INSEAD, 77305 Fontainebleau, France

The Wharton School, Department of Management, University of Pennsylvania, Philadelphia, Pennsylvania 19104

maurizio.zollo@insead.fr • winter@wharton.upenn.edu

Abstract

This paper investigates the mechanisms through which organizations develop dynamic capabilities, defined as routinized activities directed to the development and adaptation of operating routines. It addresses the role of (1) experience accumulation, (2) knowledge articulation, and (3) knowledge codification processes in the evolution of dynamic, as well as operational, routines. The argument is made that dynamic capabilities are shaped by the coevolution of these learning mechanisms. At any point in time, firms adopt a mix of learning behaviors constituted by a semiautomatic accumulation of experience and by deliberate investments in knowledge articulation and codification activities. The relative effectiveness of these capability-building mechanisms is analyzed here as contingent upon selected features of the task to be learned, such as its frequency, homogeneity, and degree of causal ambiguity. Testable hypotheses about these effects are derived. Somewhat counterintuitive implications of the analysis include the relatively superior effectiveness of highly deliberate learning processes such as knowledge codification at lower levels of frequency and homogeneity of the organizational task, in contrast with common managerial practice.

(Organizational Learning; Dynamic Capabilities; Organizational Routines; Knowledge Codification; Knowledge Articulation; Learning by Doing; Task Frequency; Task Heterogeneity; Causal Ambiguity)

Introduction

Explaining the variation in the degree of success of business organizations by reference to different degrees and qualities of organizational knowledge and competence has been a major focus of recent theorizing in both strategic management and organizational theory. Concepts and labels coined to characterize the phenomenon abound. From pioneering efforts such as Selznick's (1957) "distinctive competence," to the more recent and refined notions of organizational routines (Nelson and

Winter 1982), absorptive capacity (Cohen and Levinthal 1990), architectural knowledge (Henderson and Clark 1990), combinative capabilities (Kogut and Zander 1992) and, finally, dynamic capabilities (Teece et al. 1997), there are decades of investment in sorting out the traits and the boundaries of the phenomena. Recent contributions (Eisenhardt and Martin 2000, Dosi et al. 2000) aim at clarifying distinctions among the various constructs. As the field progresses in the characterization of the phenomena, however, the need for a better understanding of the origins of capabilities becomes increasingly apparent. What ultimately accounts for the fact that one organization exhibits "competence" in some area, while another does not? And how do we explain the growth and decay of that particular competence, other than the simple repetition, or lack thereof, of decisions and behavior?

This paper sets forth a theoretical account of the genesis and evolution of dynamic capabilities that focuses on the role of several learning mechanisms and their interaction with selected attributes of the organizational task being learned. We begin our discussion by presenting a general framework linking learning mechanisms to the evolution of dynamic capabilities, and these to the evolution of operating routines. We then draw on evolutionary economics to develop a more fine-grained representation of how organizational knowledge evolves at the level of operating routines as well as of dynamic capabilities. In §3, we discuss the different learning mechanisms in terms of a spectrum of learning investments and identify the contextual features that shape their roles. Finally, in §4, we explore in detail how some key attributes of the organizational task under consideration affect the relative effectiveness of these learning mechanisms and develop testable hypotheses for future empirical work.

1. Organizational Learning and Dynamic Capabilities

Our first objective is to develop a framework describing the linkages among the processes we intend to study. We

describe a set of learning mechanisms encompassing both the relatively passive experiential processes of learning (“by doing”) and more deliberate cognitive processes having to do with the articulation and codification of collective knowledge. These learning processes are responsible for the evolution in time of two sets of organizational activities: one geared towards the operational functioning of the firm (both staff and line activities), which we will refer to as *operating routines*; the other dedicated to the modification of operating routines, which we identify with the notion of *dynamic capabilities*.

Below, we provide a more detailed description of how we conceive of dynamic capabilities and of the function of the identified learning mechanisms.

Dynamic Capabilities

Building on Nelson and Winter’s view of the organization as a set of interdependent operational and administrative routines which slowly evolve on the basis of performance feedbacks, Teece et al. (1997) define the concept of “dynamic capabilities” as “the firm’s ability to integrate, build, and reconfigure internal and external competencies to address rapidly changing environments” (p. 516). While this suggests something of what dynamic capabilities are for and how they work, it leaves open the question of where they come from. Also, the definition seems to require the presence of “rapidly changing environments” for the existence of dynamic capabilities, but firms obviously do integrate, build, and reconfigure their competencies even in environments subject to lower rates of change. We propose the following alternative:

DEFINITION. A dynamic capability is a learned and stable pattern of collective activity through which the organization systematically generates and modifies its operating routines in pursuit of improved effectiveness.

In addition to avoiding the near tautology of defining capability as ability, this definition has the advantage of specifically identifying operating routines, as opposed to the more generic “competencies,” as the object on which dynamic capabilities operate. Also, it begins to spell out some of the characteristics of this construct. The words “learned and stable pattern” and “systematically” highlight the point that dynamic capabilities are structured and persistent; an organization that adapts in a creative but disjointed way to a succession of crises is not exercising a dynamic capability. Dynamic capability is exemplified by an organization that adapts its operating processes through a relatively stable activity dedicated to process improvements. Another example is given by an organization that develops from its initial experiences with acquisitions or joint ventures a process to manage such projects in a systematic and relatively predictable fashion.

The ability to plan and effectively execute postacquisition integration processes is another example of a dynamic capability, as it involves the modification of operating routines in both the acquired and the acquiring unit.

Figure 1, below, offers a generic picture of the organizational processes under consideration and of the linkages between them. Dynamic capabilities arise from learning; they constitute the firm’s systematic methods for modifying operating routines. To the extent that the learning mechanisms are themselves systematic, they could (following Collis 1994) be regarded as “second-order” dynamic capabilities. Learning mechanisms shape operating routines directly as well as by the intermediate step of dynamic capabilities.

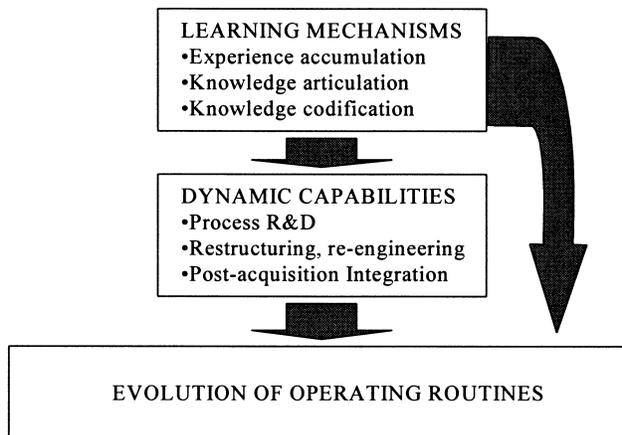
Learning Mechanisms

What mechanisms are involved in the creation and evolution of dynamic capabilities? What features distinguish an organization capable of systematically developing new and enhanced understanding of the causal linkages between the actions it takes and the performance outcomes it obtains?

In the spirit of bridging the behavioral and cognitive approaches to the organizational learning phenomenon (Glynn et al. 1994), we attend both to the experience accumulation process and to more deliberate cognitive processes involving the articulation and codification of knowledge derived from reflection upon past experiences.¹ The three mechanisms focal to this analysis are introduced and discussed separately below.

Organizational Routines and Experience Accumulation. Routines are stable patterns of behavior that characterize organizational reactions to variegated, internal or external stimuli. Every time an order is received from a

Figure 1 Learning, Dynamic Capabilities, and Operating Routines



customer or a decision is made to upgrade a production process, for instance, a host of predictable and interrelated actions are initiated which will eventually conclude with the shipping of the ordered goods (and receipt of corresponding payment) or with the launch of the new production system. In spite of the superficial similarity between these two examples, though, the two patterns of behavior present a theoretically relevant distinction. The first type of routine involves the execution of known procedures for the purpose of generating current revenue and profit, while the second seeks to bring about desirable changes in the existing set of operating—in this case, production—routines for the purpose of enhancing profit in the future. Routines of the second type are traditionally identified as *search* routines in evolutionary economics (Nelson and Winter 1982), and are here regarded as constitutive of dynamic capabilities.

Routines of the two types have different effects on the generation and appropriation of rents, depending on the pace of change in the environment. Of course, effective operating routines are always a necessity, and superior operating routines are always a source of advantage. In a relatively static environment, a single learning episode may suffice to endow an organization with operating routines that are adequate, or even a source of advantage, for an extended period. Incremental improvements can be accomplished through the tacit accumulation of experience and sporadic acts of creativity. Dynamic capabilities are unnecessary, and if developed may prove too costly to maintain. But in a context where technological, regulatory, and competitive conditions are subject to rapid change, persistence in the same operating routines quickly becomes hazardous. Systematic change efforts are needed to track the environmental change; both superiority and viability will prove transient for an organization that has no dynamic capabilities. Such capabilities must themselves be developed through learning. If change is not only rapid but also unpredictable and variable in direction, dynamic capabilities and even the higher-order learning approaches will themselves need to be updated repeatedly. Failure to do so turns core competencies into core rigidities (Leonard Barton 1992).

To our knowledge at least, the literature does not contain any attempt at a straightforward answer to the question of how routines—much less dynamic capabilities—are generated and evolve. That is hardly surprising, for it seems clear that routines often take shape at the conjunction of causal processes of diverse kinds—for example, engineering design, skill development, habit formation, particularities of context, negotiation, coincidences, efforts at imitation, and ambitions for control confronting

aspirations for autonomy. In general, however, the literature that is relatively explicit on the question suggests that “routines reflect experiential wisdom in that they are the outcome of trial and error learning and the selection and retention of past behaviors” (Gavetti and Levinthal 2000, p. 113). This view is linked to an emphasis on the importance of tacit knowledge, since tacitness arises when learning is experiential. It is also linked to the notion that organizational routines are stored as procedural memory (Cohen and Bacdayan 1994) and to the image of routinized response as quasi-automatic. In addition, it is consistent with the traditional view of organizational learning as skill building based on repeated execution of similar tasks that is implicit in much of the empirical literature on learning curves (e.g., Argote 1999).

We incorporate this view in our discussion here, using the term “experience accumulation” to refer to the central learning process by which operating routines have traditionally been thought to develop. In adopting this usage, we seek to provide a clear baseline for our analysis of some of the other learning processes that shape routines and dynamic capabilities. With respect to the latter, it seems clear that a theory of their development and evolution must invoke mechanisms that go beyond semi-automatic stimulus-response processes and tacit accumulation of experience.

Knowledge Articulation. One of the recognized limitations of the behavioral tradition in the study of organizational learning consists of the lack of appreciation of the deliberative process through which individuals and groups figure out what works and what doesn’t in the execution of a certain organizational task (Cangelosi and Dill 1965, p. 196; Levinthal and March 1988, p. 208; Narduzzo et al. 2000). Important collective learning happens when individuals express their opinions and beliefs, engage in constructive confrontations and challenge each other’s viewpoints (Argyris and Schon 1978, Duncan and Weiss 1979). Organizational competence improves as members of an organization become more aware of the overall performance implications of their actions, and is the direct consequence of a cognitive effort more or less explicitly directed at enhancing their understanding of these causal links. We therefore direct attention to a second mechanism of development of collective competence, the process through which implicit knowledge is articulated through collective discussions, debriefing sessions, and performance evaluation processes. By sharing their individual experiences and comparing their opinions with those of their colleagues, organization members can achieve an improved level of understanding of the causal mechanisms intervening between the actions required to

execute a certain task and the performance outcomes produced. Organizational processes are often subject to significant causal ambiguity with respect to their performance implications (Lippman and Rumelt 1982), and particularly so in rapidly changing environmental contexts. Higher-level cognitive efforts and a more deliberate collective focus on the learning challenge can help to penetrate the ambiguity—although some part of it always persists. It is important to note that only a small fraction of articulable knowledge is actually articulated, and that organizations differ substantially on the degree to which they transform potentially articulable knowledge into articulated statements (Winter 1987, Kogut and Zander 1992, Cowan, et al. 2000). While potentially requiring significant efforts and commitment on the part of the members of the organization, such articulation efforts can produce an improved understanding of the new and changing action-performance links, and therefore result in adaptive adjustments to the existing sets of routines or in enhanced recognition of the need for more fundamental change.

Knowledge Codification. An even higher level of cognitive effort is required when individuals codify their understandings of the performance implications of internal routines in written tools, such as manuals, blueprints, spreadsheets, decision support systems, project management software, etc. Knowledge codification is a step beyond knowledge articulation. The latter is required in order to achieve the former, while the opposite is obviously not true. The fact that in most cases articulated knowledge is never codified bears witness to additional costs incurred when stepping up the learning effort from a simple sharing of individual experience to developing manuals and other process-specific tools.

Note that while some of these tools are deliberately aimed at uncovering the linkages between actions and performance outcomes (such as performance appraisals, postmortem audits, etc.), most of them are intended simply to provide guidelines for the execution of future tasks. Whatever the intentions motivating the codification effort, the process through which these tools are created and consistently updated implies an effort to understand the causal links between the decisions to be made and the performance outcomes to be expected, even though learning might not be the deliberate goal of the codification effort. To illustrate, both the readers and the authors of this article have certainly experienced the significant increase in the clarity of their ideas consequent to the act of writing the first draft of a research paper. Through the writing process, one is forced to expose the logical steps of one's arguments, to unearth the hidden assumptions,

and to make the causal linkages explicit. Similarly, a group of individuals who are in the process of writing a manual or a set of written guidelines to improve the execution of a complex task (think of the development of a new product, or the management of the postacquisition integration process) will most likely reach a significantly higher degree of understanding of what makes a certain process succeed or fail, compared to simply telling "war stories" or discussing it in a debriefing session.

Knowledge codification is, in our view, an important and relatively underemphasized element in the capability-building picture. The literature has emphasized that codification facilitates the diffusion of existing knowledge (Winter 1987, Zander and Kogut 1995, Nonaka 1994), as well as the coordination and implementation of complex activities. Having identified and selected the change in the operating routines or the new routine to be established, the organization should create a manual or a tool to facilitate its replication and diffusion. The principal benefits to the codification effort are seen as coming from the successful use of the manual or tool.

While this picture is clearly accurate in many cases, it may also exaggerate some benefits of codification while neglecting others. It overlooks, for example, important obstacles to the success of this type of use of codified knowledge, such as the difficulty of assuring that the codified guidance is both adequate and actually implemented. More importantly, this appraisal overlooks the learning benefits of the activities necessary for the creation and development of the tools. To develop a manual for the execution of a complex task, the individuals involved in the process need to form a mental model of what actions are to be selected under what conditions. By going through that effort, they will most likely emerge with a crisper definition of what works, what doesn't work, and why.

Codification, therefore, is potentially important as a supporting mechanism for the entire knowledge evolution process, not just the transfer phase. It can, for instance, facilitate the generation of new proposals to change the currently available routines, as well as the identification of the strengths and the weaknesses in the proposed variations to the current set of routines. The cognitive simplification inherent in the act of synthesizing on paper (or in a computer program) the logic behind a set of instructions can therefore represent both an economizing on data-processing requirements, allowing more effective decisionmaking (Boisot 1998, Gavetti and Levinthal 2000), and a more-or-less deliberate form of retrospective sense-making with respect to the performance implications of a given set of activities (Weick 1979 and 1995). This sense-making typically draws upon a specialized

language and causal understanding developed for and by other codification efforts of a similar type (Cowan and Foray 1997, Cowan et al. 2000).

To be sure, these advantages do not come for free. There are specific costs attached to the knowledge codification process. Direct costs include the time, the resources, and the managerial attention to be invested in the development and updating of task-specific tools, while indirect costs include a possible increase in the rate of “misfire” or inappropriate application of the routine (Cohen and Bacdayan 1994) if the codification is poorly performed, and the more general increase in organizational inertia consequent to the formalization and structuring of the task execution. Concern with these costs is legitimate and familiar. A long debate, dating back to Weber’s work, has engaged organizational theorists with respect to the advantages and disadvantages of formalization, a kindred phenomenon to knowledge codification. For a long time, the skeptics seemed to have the upper hand. More recently, there seems to be increasing willingness to see the formalization of operating routines in a positive light; e.g., as sometimes capable of producing an “enabling” rather than a “coercive” bureaucracy (Adler and Borys 1996).

This contingency approach is consistent with the thrust of the current analysis: Instead of rejecting knowledge codification processes *tout court* as producers of inertial forces, our objective is to determine the conditions under which the learning and diffusion advantages attached to codification could more than offset its costs. We acknowledge, of course, that this is not likely to be true in all cases—and that codification, like many other things, is likely to produce bad results when done badly. In the concluding section, we discuss what seems to be involved in doing it well.

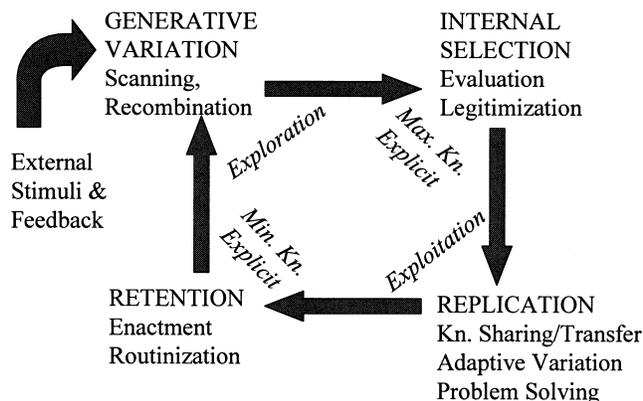
2. The Cyclical Evolution of Organizational Knowledge

Having offered our view on the notion of dynamic capabilities and the learning processes that might support their evolution, we now provide a more detailed description of the way in which dynamic capabilities and operating routines evolve in time.

Figure 2 offers a graphical view of a “knowledge evolution cycle.” It is a simple account of the development of collective understanding regarding the execution of a given organizational task. We have adapted for our purposes the classic evolutionary paradigm of variation-selection-retention.²

Organizational knowledge is here described as evolving through a series of stages chained in a recursive cycle.

Figure 2 Activities in the Knowledge Evolution Cycle



The logical point of departure for each cycle lies in the variation stage, where individuals or groups of them generate a set of ideas on how to approach old problems in novel ways or to tackle relatively new challenges. This happens on the basis of a combination of external stimuli (competitors’ initiatives, normative changes, scientific discoveries, etc.) with internally generated information derived from the organization’s existing routines, and may involve substantial creativity. These sets of ideas, initially in embryonic and partly tacit form, are then subject to internal selection pressures aimed at the evaluation of their potential for enhancing the effectiveness of existing routines or the opportunity to form new ones (Nonaka 1994). These pressures arise as new ideas are considered in relation to a shared understanding of the organization’s prior experience, as well as in the context of established power structures and existing legitimization processes. Hypotheses as to the expected advantages from the application of the proposed changes are probed through a collective investment in articulation, analysis, and debate of the merits and risks connected to the initiative.

The third phase of the cycle refers to the set of activities enacted by an organization for the purpose of diffusing the newly approved change initiatives to the relevant parties within the firm. This diffusion process requires the spatial replication of the novel solutions in order to leverage the newly found wisdom in additional competitive contexts (see Winter and Szulanski 2001). This mechanism is a relatively new addition to the standard variation-selection-retention triumvirate of evolutionary modeling borrowed from evolutionary biology (Campbell 1969). Its rationale lies in the simple observation that organizations differ in a crucial way from biological entities in that they act simultaneously in spatially diverse contexts, posing the question of how novelties (ideas or change proposals

in our case) that survive internal selection processes diffuse to the spatially dispersed parts of the organizations where the novel approach can be put to use.³

The replication phase, however, does not simply serve the function of diffusing the newly minted knowledge to the places and times where it is needed. Importantly for the purposes of this paper, the diffusion of knowledge through replication processes also contributes new (raw) information that can provide the diversity needed to start a variation phase of a new knowledge cycle. The application of the routines in diverse contexts generates new information as to the performance implications of the routines employed. The hypotheses constructed through the cognitive efforts of the selection phase can now (in theory) be tested with empirical evidence. This evidence, if properly collected and processed, can further illuminate the context-dependent cause-effect linkages between decisions and performance outcomes. It can thereby prime the initiation of a new knowledge cycle.

The qualification concerning collection and processing is important because replication and repetition processes tend to make knowledge evolve toward a more tacit form as it becomes highly embedded in the behavior of the individuals involved in the multiple executions of the task. Repetition leads to automaticity in the execution of a given task and to a corresponding reduction in individual awareness and collective understanding of the action-performance linkages, as well as of the purpose of the execution criteria followed. The abundance of raw information potentially available due to the repeated execution of the new tasks is, in fact, typically unusable without significant cognitive efforts aimed at evaluating, classifying, analyzing, and finally distilling the results into new initiatives to modify existing routines or create new ones.

The external environment plays two distinguishable roles in the process. It supplies diverse stimuli and substance for internal reflections on possible applications to the improvement of existing routines. Of course, it also functions as a selection mechanism in the classic evolutionary sense as it provides the feedback on the value and viability of the organization's current behaviors. While we fully recognize the fundamental relevance of both these roles, the focus of this paper is on the set of internal processes located, from a temporal standpoint, between the two.

A few more observations can be made by entering some additional elements in the basic framework. First, the knowledge cycle proceeds, in March's (1991) terms, from an exploration phase to an exploitation one, potentially feeding back into a new exploration phase. Exploration activities are primarily carried out through cognitive efforts aimed at generating the necessary range of

new intuitions and ideas (variation) as well as selecting the most appropriate ones through evaluation and legitimization processes. By contrast, exploitation activities rely more on behavioral mechanisms encompassing the replication of the new approaches in diverse contexts and their absorption into the existing sets of routines for the execution of that particular task. We suggest that exploitation can prime exploration, and propose that in addition to the familiar trade-off between exploration and exploitation processes, there can be a recursive and co-evolutionary relationship between them. This may indicate a way to conceptualize the managerial challenge of handling both processes simultaneously.

The second observation is that the nature of organizational knowledge changes over the cycle. In the first phases of generative variation and internal selection, the initial idea or novel insight needs to be made increasingly explicit to allow a debate on its merits, and the knowledge reaches a peak of explicitness as the cycle reaches the selection stage. Through the replication and retention phases then, knowledge becomes increasingly embedded in human behavior and likely gains in effectiveness while declining in abstraction (as it is applied to a wider variety of local situations) and in explicitness (as even the people involved in its application have difficulty in explaining what they are doing and why). The changing nature of knowledge throughout the evolutionary cycle is an issue of primary concern to us as well as to some other scholars (Nonaka 1994, Nonaka and Takeuchi 1995, Boisot 1998). We here identify as sources of such variations the different degrees of intentionality necessary at different stages of the cycle, and construct a theoretical argument about the conditions that warrant increasing efforts to turn implicit into explicit knowledge.

3. Learning Investments and Dynamic Capability Building

We now concentrate our analysis on the complex interplay of the learning mechanisms with the change processes that come under the heading of dynamic capabilities. The following proposition summarizes our view of the dynamic capability-building process:

PROPOSITION. *Dynamic capabilities emerge from the coevolution of tacit experience accumulation processes with explicit knowledge articulation and codification activities.*

The reference to coevolution signals the importance of viewing the influence of the three mechanisms in terms of continuing interaction and mutual adjustment. Our emphasis is, however, not so much on the causal links among

the three capability-building mechanisms as on the way their interaction affects the firm's ability to improve the existing set of operating routines. In other words, firms learn systematic ways to shape their routines by adopting an opportune mix of behavioral and cognitive processes, by learning how to articulate and codify knowledge, while at the same time they facilitate the accumulation and absorption of experiential wisdom. The ambition here is to begin an analytical effort aimed at providing a clearer sense of what an "opportune" mix of learning processes might mean. In this section, we frame the relationships among the mechanisms in terms of varying levels of investment in deliberate learning activities. We then attempt to identify the costs connected with these activities. Finally, we identify some of the most important contingencies that need to be considered to understand how costs and benefits vary in different contexts.

Learning Investments. To study the relationships among the three learning mechanisms, we consider their positions on a continuum in the investment of resources (financial, temporal, and cognitive) directed towards the purpose of improving the collective understanding of action-performance linkages.

The level of investment in developing dynamic capabilities will be the lowest when the firm counts on the experience accumulation process, as the learning happens in an essentially semiautomatic fashion on the basis of the adaptations that individuals enact in reaction to unsatisfactory performance. The effectiveness of this learning mechanism requires "only" the stability of personnel exposed to the experienced events, good performance-monitoring systems, and sufficiently powerful incentives to ensure that individuals will initiate the search routines when performance levels decay. Examples of activities that rely primarily on this type of learning mechanism include the creation of a specific function or department responsible for the process to be learned, and the hiring of "specialists" in the execution of the task under scrutiny (e.g., creating an M&A team, hiring a TQM expert, defining the function of the chief knowledge officer). These actions can be understood as efforts to provide a specific organizational locus where experience can accumulate.

The "learning investment" is likely to be higher when the organization (or the relevant unit) relies on knowledge articulation processes to attempt to master or improve a certain activity. In addition to the actions mentioned above as necessary to facilitate experience accumulation processes, the organization will have to incur costs due to the time and energy required for people to meet and discuss their respective experiences and beliefs (Ocasio 1997). A meeting organized to debrief the participants in

a complex project can be expensive in terms of both the direct costs and, most importantly, the opportunity costs deriving from the sacrifice of time dedicated to working on active projects. Opportunity cost considerations have the somewhat paradoxical effect of tending to suppress learning when it is most valuable and needed: The higher the activity levels in the execution of a certain task, the higher the opportunity costs for the learning investments dedicated to that specific task, and therefore the lower the likelihood that the hyperactive team will afford the time to debrief, despite the obvious advantages from the potential identification of process improvements (Tyre and Orlikowski 1994).

The investment of time, efforts, and resources will be the highest in the case of knowledge codification processes. Here, the team involved in the execution of the task not only has to meet and discuss, but also has to actually develop a document or a tool aimed at the distillation of the insights achieved during the discussion(s). If a tool (a manual or a piece of software) already exists, the team has to decide whether and how to update it, and then do it.

It is clear that, given a certain task, organizations might differ quite substantially in the degree to which they decide to invest in these kinds of deliberate, cognitively intense learning activities. Consider by way of example the approaches taken by Hewlett Packard and Corning in developing their internal competence in the management of strategic alliances.⁴ Both companies are considered to be among the most experienced and sophisticated players in the strategic alliance arena. The ways in which they managed their capability development processes, though, were at the opposite extremes of our learning investment continuum. Whereas Corning made a point of avoiding the codification of its alliance practices, preferring to rely on experience-based learning and apprenticeship systems, HP decided to invest in a large variety of collective learning mechanisms to both identify and diffuse the best practices in this particular task. They started with the creation of a database of all their past alliances, for each of which they requested a written postmortem analysis. They then proceeded to create what became a 400-page binder with the distilled wisdom on the criteria to follow and the pitfalls to avoid in each phase of the alliance process. Further, the expert team in charge of the alliance practice initiated a series of internal workshops/seminars open to any HP manager to ensure the diffusion of their alliance management practices. In the language of our framework, HP combined support of behavioral learning mechanisms (the construction of a stable team of alliance experts) with deliberate investments in knowledge articulation (regular

debriefing sessions) and codification (written postmortem analysis and updating of the alliance manual).

In terms of learning investments, HP's commitment of managerial resources is much higher than Corning's. But was the investment worthwhile? What kind of returns can organizations reap from higher levels of commitment to deliberate learning processes? And, most importantly, what factors influence those returns? What weight should be given to the fact that the Corning approach, relying heavily on tacit accumulation, is more vulnerable to personnel turnover? Ex post, both approaches seem successful, suggesting some degree of equifinality, but which one is preferable, ex ante, for a firm that is trying to develop a similar competence?

To start answering these questions, we need to examine some of the contingencies under which higher degrees of cognitive effort are likely to be justifiable in terms of greater learning effectiveness. Three broad categories of factors are particularly relevant to our purposes:

Environmental Conditions (such as the speed of technological development or the time-to-market lags required by customers). Consider the high-speed end: Are investments in knowledge articulation and codification more justifiable in the context of high-tech industries? A trade-off is relevant here. On the one hand, the cognitive simplification afforded by knowledge codification has been argued to be advantageous and preferable to behavioral adaptations in these types of industries (Brown and Eisenhardt 1997). On the other, speed requirements in operating routines raise the opportunity costs of deliberate learning investments, particularly codification. In less turbulent environments, learning processes based on craftsmanship, as in Brown and Duguid's (1991) notion of communities of practice, appear to be both more effective and cheaper than their highly inertial alternatives.

Organizational Features. Organizations differ in their dynamic capabilities partly because they inhabit environments with differing rates of change, but also partly because they place different bets, implicitly or explicitly, on the strategic importance of change in the future. Organizations that are culturally disposed toward betting on change, or whose managements have successfully instilled an acceptance of continual change practices, are likely to obtain higher returns from learning at any given level of learning investment because they are more effective in shifting behavior to exploit the novel understandings. From a design standpoint, the presence of structural barriers among the different activities, typical, for example, of multidivisional firms, will raise the returns from deliberate learning investments, since experiential processes will tend to remain localized to the source of novel

insights unless mobilized through explicit learning processes. Finally, scale economies in the use of information are certainly relevant, as the standard view of codification emphasizes. Hence, larger, more divisionized and diversified organizations, as well as firms with stronger change-prone cultures, are likely to have greater opportunities to benefit from deliberate learning investments.

Task Features. Returns to investments in deliberate learning mechanisms are likely to be influenced also by a number of dimensions characterizing the task that the organization is trying to master. A task of higher economic importance, for example, will clearly justify a relatively higher investment in cognitive activities aimed at developing competence and avoiding future failures. Another relevant dimension is the scope of the task, which determines whether a group, a department, or an entire organization is responsible for its execution. The larger the scope, the more compelling the recourse to more deliberate learning mechanisms, since individuals are able to experience, and therefore understand, only a small fraction of the task.

The implications of some other task features are less obvious, however. In the following section, we examine in detail three determinants of the returns to learning investments. We see substantial potential in the study of how knowledge accumulation, articulation, and codification processes interact with task features. First, the analysis is at a level amenable to strategic action on the part of the firm (while there is relatively little the firm can do to operate on its own cultural features or change its environmental context). Second, this emphasis tends to correct a bias that has arguably distorted the organizational learning literature: The organization is typically considered to be engaged in highly frequent, relatively homogeneous types of tasks, reasonably well defined in their decisional alternatives and sometimes even in their action/performance linkages. This is the case with the vast majority of the learning studies conducted so far in organizational settings, where the task typically observed consists of manufacturing or service operations.

4. The Moderating Role of Task Features

The central tenet of the present section is that the relative effectiveness of the learning mechanisms depends on the characteristics of the tasks that the organization is attempting to learn and of the operating routines that it is interested in adjusting or radically redesigning.

We will focus our attention on three specific dimensions of the task or operating routines at hand. The first

has to do with its frequency, or how often it gets triggered and executed within a specific period of time. The second is its degree of heterogeneity, or how novel the task appears each time to the unit that has to execute it. It is a matter of the dispersion in the defining traits of the task across multiple occurrences. The third dimension relates to the degree of causal ambiguity in the action-performance links, or how easy it is to derive clear indications as to what should or should not be done in the execution of the task. We will elaborate on each of these in turn.

Frequency. It is clear that organizational tasks vary immensely on this dimension. In principle all of them, from extremely high frequencies such as check processing in a bank to rare events such as a reorganization or a CEO search process, are subject to learning mechanisms. The question of what mechanism(s) work better at different frequency levels has been substantially neglected, though March et al. (1991) have insightfully explored the situation at the very bottom of the spectrum. Based on the model presented above, we argue that at increasing frequency levels, the capability-building mechanism based on tacit accumulation of experiences in the minds of “expert” personnel becomes increasingly effective as a learning mechanism relative to the more explicit investments in knowledge articulation and knowledge codification processes. At lower frequency levels, while all the mechanisms will suffer significant losses in their capability-building power, we argue that the relative effectiveness ranking inverts, and knowledge codification becomes increasingly effective relative to knowledge articulation, which in turn becomes more effective than tacit experience accumulation. The reasons are as follows.

Individual Memory. The experience accumulation mechanism relies on the memory of individuals exposed to previous occurrences. Other things being equal, this suggests the more frequent the event is, the higher the likelihood that individuals will have retained their impressions as to what worked and what didn’t work in the previous experiences. Indeed, the success of codification efforts may be limited because the results of tacit learning are too entrenched and people might believe the consultation and application of task-specific tools to be redundant.

Coordination Costs. The knowledge articulation and codification mechanisms become increasingly complex and costly to coordinate as the frequency of the event increases. Individuals need to meet to brainstorm, and typically need face-to-face contact to coordinate the completion or upgrading of a manual or a decision-support system.

Opportunity Costs. Conducting debriefing sessions and updating tools after the completion of the task cannot be done too often without diverting attention away from day-to-day operations. A balance between explicit learning activities and execution activities, between thinking and doing, is essential (March 1991, Mukherjee et al. 1999). March et al. (1991) argue that with highly infrequent events, organizations can learn from quasihistories (i.e., “nearly happened” events) or from scenario analysis. Both mechanisms entail a substantial amount of investment in cognitive efforts and, most likely, rely on the creation of written output or on the use of electronic support systems to identify and make all the assumptions explicit.

These arguments can be expressed in a slightly more formal way by advancing the following hypothesis.

HYPOTHESIS 1. *The lower the frequency of experiences, the higher the likelihood that explicit articulation and codification mechanisms will exhibit stronger effectiveness in developing dynamic capabilities, as compared with tacit accumulation of past experiences.*

The comparison of the results obtained by a few recent studies on the effectiveness of learning processes in relatively low-frequency tasks seems to lend some support to this contingency argument. As task frequency increases, moving from reengineering processes (Walston 1998) to acquisitions (Zollo 1998) to alliances (Kale 1999), and finally to quality improvement projects (Mukherjee et al. 1999), the relative effectiveness of experience accumulation with respect to knowledge articulation and codification processes seems to increase. Whereas in reengineering processes prior experiences do not affect the performance of the current project, while efforts to create written guidelines do (Walston 1998), in the case of relatively more frequent quality improvement efforts, cognitive (modeling and simulation processes) and experiential (i.e., experiments) learning mechanisms show comparable (and statistically significant) effectiveness (Mukherjee et al. 1999). Acquisitions and alliances typically occur at intermediate frequency levels compared to reengineering and quality improvement processes. Kale (1999) and Zollo (1998) show consistent results pointing to a weak but positive role for the experience accumulation mechanism and a stronger performance effect for the knowledge codification proxy (number of task-specific tools developed by the focal firm).⁵

Heterogeneity. The variance in the characteristics of the task as it presents itself in different occurrences presents a different, albeit related, type of challenge with respect to the frequency problem. The issue here is that individuals have to make inferences as to the applicability

of lessons learned in the context of past experiences to the task presently at hand. As task heterogeneity increases, inferences become more difficult to make and, when made, they are more likely to generate inappropriate generalizations and poorer performance (Cormier and Hagman 1987, Gick and Holyoak 1987, Holland et al. 1986, Holyoak and Thagard 1995). How does the degree of task heterogeneity affect the relative effectiveness of the capability-building mechanisms? We submit that the more explicit mechanisms will be relatively more effective in developing dynamic capabilities, compared to tacit experience accumulation, at higher degrees of task heterogeneity. The rationale is that the hazards of inappropriate generalization can only be attenuated via an explicit cognitive effort aimed at uncovering the interdependence between the dimension(s) of heterogeneity and the action-performance relationships. For example, a firm that has made several acquisitions in a wide variety of sectors will probably find it more difficult to extrapolate rules of conduct in managing acquisition processes, compared to another one that has consistently acquired in its own domain. The former might find it comparatively more useful to invest in debriefing sessions and in detailed postmortem analyses as opposed to simply relying on its group of M&A experts. The need to understand what works and what doesn't in the different contexts experienced requires an explicit investment in retrospective sense-making, which fosters the development of specific capabilities to address the different contexts in which acquisitions might be completed in the future.

Empirical evidence unearthed in some recent studies of acquisition performance seems to support this line of reasoning. Haleblan and Finkelstein (1999) find a U-shaped learning curve connecting prior acquisition experience with the performance of the focal acquisition, whereas Zollo and Reuer (2000) replicate and extend this finding to the impact of prior alliance experience on the performance of the focal acquisition. The initially counterintuitive result represented by the declining portion of the U curve is plausibly interpreted as a consequence of the high heterogeneity in this kind of organizational task, which increases the likelihood of erroneous generalizations. Consistent with this interpretation, knowledge codification processes have been shown to be strongly related to performance in these conditions, and therefore relatively more effective than simple experience accumulation processes (Zollo 1998)

These arguments suggest the following hypothesis for future empirical work.

HYPOTHESIS 2. *The higher the heterogeneity of task experiences, the higher the likelihood that explicit articulation and codification mechanisms will exhibit stronger effectiveness in developing dynamic capabilities, as compared with tacit accumulation.*

ulation and codification mechanisms will exhibit stronger effectiveness in developing dynamic capabilities, as compared with tacit accumulation.

Causal Ambiguity. The third task dimension that we take into consideration is the level of causal ambiguity, or (conversely) the degree of clarity in the causal relationships between the decisions or actions taken and the performance outcomes obtained (Lippman and Rumelt 1982). Irrespective of the degree of expertise developed in handling a certain task, there are a number of factors that obscure these cause-effect linkages. The number and the degree of interdependence of subtasks obviously are important considerations affecting the uncertainty as to the performance implications of specific actions. These, in fact, are the two key parameters of complexity addressed in the "NK" formal modeling approach, and it is complexity that poses the obstacle to learning by local "experiential" search (Gavetti and Levinthal 2000). Another important factor is the degree of simultaneity among the subtasks. If the subtasks can be managed in a sequential fashion, it will be easier to pinpoint the consequences of each part for the performance of the entire process.

Again, the costs related to the "learning investments" described above will be justified and justifiable in the presence of high causal ambiguity, as the higher degrees of cognitive effort implicit in the articulation and codification of the lessons learned in previous experiences should help penetrate the veil of ambiguity and facilitate the adjustment of the routines.

We therefore submit this final hypothesis for testing in future empirical work:

HYPOTHESIS 3. *The higher the degree of causal ambiguity between the actions and the performance outcomes of the task, the higher the likelihood that explicit articulation and codification mechanisms will exhibit stronger effectiveness in developing dynamic capabilities, as compared with tacit accumulation of past experiences.*

5. Conclusions

This paper proposes a coherent structure for the study of the formation and evolution of dynamic capabilities within organizations. It does so by drawing on arguments derived from both the behavioral and cognitive traditions in organizational learning studies. Starting from a characterization of dynamic capabilities as systematic patterns of organizational activity aimed at the generation and adaptation of operating routines, we have proposed that they develop through the coevolution of three mechanisms: tacit accumulation of past experience, knowledge articulation, and knowledge codification processes.

One implication of the analysis is the perhaps provocative suggestion that knowledge codification (and to a lesser extent knowledge articulation) activities become superior mechanisms with respect to the accumulation of expertise as the frequency and the homogeneity of the tasks are reduced. This learning-oriented appraisal runs counter to the logic of codification that now dominates both theory and practice. Ordinarily, a bank would copiously codify its branch operations (how to open an account, execute a wire transfer, etc.) and a manufacturer would do the same with its standard operating procedures, but neither would typically be prepared to do so when it comes to managing a reengineering process or the acquisition of a company. This is the natural result of a habit of thinking that the costs of codification activities are justified only by their outputs, and not by the learning benefits of the codification process itself. In our perspective, the creation of a manual or a decision support tool aimed at the facilitation of a relatively infrequent and heterogeneous task may be more valid (or at least as valid) as a capability-building exercise as for the benefit derived from the actual use of the tool. It can affect the level of performance on subsequent tasks, even if limited in number, and it can affect performance even if individual minds, and not the finished tool, are the key repository of the improved understanding. Codification efforts force the drawing of explicit conclusions about the action implications of experience, something that articulation alone (much less experience alone) does not do.

While the standard logic of codification correctly directs attention to the costs of codification and the scale of application, the learning perspective suggests other important desiderata for codification to be “done right.” Applying the logic of our argument here, and drawing on a considerable amount of related empirical work and discussion,⁶ we propose four guiding principles. First, codification should aim at developing and transferring “know why” as well as “know how.” We have emphasized that codification efforts provide an occasion for valuable efforts to expose action-performance links. Aiming at process prescriptions alone forfeits this advantage and increases risks of inappropriate application. Second, codification efforts should be emphasized at an appropriate time in the course of learning. Attempted prematurely, codification efforts risk hasty generalization from limited experience, with attendant risks of inflexibility and negative transfer of learning. When codification is long deferred, attempts at careful examination of the causal relations may be frustrated by the entrenched results of tacit accumulation, which may have attained high acceptance for both superstitious (Levitt and March 1988) and ra-

tional (March and Levinthal 1993) reasons. What constitutes the “right time” depends on the overall heterogeneity of the tasks and the representativeness of the early sample; high heterogeneity and low representativeness urge deferral.

The third principle calls for the codified guidance to be tested by adherence. The implementation issues here are particularly delicate. It is clear, to begin with, that codification cannot be an instrument of continuing learning if the guidance developed after Trial 7 is simply ignored at Trial 8 and later. Subsequent application provides the experimental test for the causal understanding that the guidance is supposed to embody; the shortcomings of that understanding cannot be identified if the guidance is too easily dismissed. Because, however, it is implicit here that the guidance could be flawed, the desirability of testing by adherence is obviously qualified by the need to avoid inappropriate application. The need for judgment is inescapable, but the fourth and final point here is that there is also a need for some supporting structure. Significant departures from the guidance should not be entirely at the discretion of the task team, but should be subject to review and approval by a body that can assess the case in the light of the longer-term interest in capability building. At a minimum, such a process can serve the function of making sure that the subsequent discussion in the selection stage of the knowledge cycle is correctly informed about the path actually taken and can consider, for example, whether the departure should be adopted as part of the standard guidance. Such a review policy brings its own hazards, of course, particularly the risk that the review process may become a bottleneck. The processing of engineering change orders presents strong parallels with the organizational issues involved here.

In our view, a more nuanced assessment of knowledge codification and, more generally, of deliberate learning processes, is just one example of the potential benefits, both for theory building and management, of an inquiry into how competence is generated and evolves within an organization. The framework introduced in this paper—particularly the knowledge evolution cycle and the relationships among learning, dynamic capabilities, and operating routines—constitutes, we believe, a significant clarification of the structure of the phenomena. This inquiry is, however, still in its infancy. We know little, for example, of how the characteristics of the organizational structure and culture interact with the features of the task to be mastered in determining the relative effectiveness of the various learning behaviors. Why is it that certain

firms, with comparable levels of expertise, codify a set of activities more than others do? And under what conditions does that enable, as opposed to inhibit, performance? To what extent is intentionality necessary to produce adaptive adjustments in existing routines? The complexity of these questions, compounded by the fact that we have only recently started to converge on a parsimonious vocabulary for these concepts, is only comparable to the magnitude of the expected returns from the advancement of our knowledge on these issues.

Beyond theory building, we hope that the present paper provides useful guidance for future empirical inquiry into the role that articulation and codification processes play in creating dynamic capabilities. That is its principal purpose, and although the existing empirical base is thin, we consider that there is already good reason to believe that significant progress in that direction is quite possible. More speculatively, we hope we have contributed in a small way to a major research thrust that seems to be emerging, the effort to expand our understanding of how cognitive activity of a deliberate kind shapes organizational learning, knowledge, and action. Although there have long been important voices urging a different course (Weick 1979, 1995), too much of our theoretical understanding of organizations and competitive processes has been framed either by the unrealistic models of hyper-rationality favored in economics or by the more realistic but still distorted versions of bounded rationality favored in the behavioral tradition. People acting in organizational contexts often think imperfectly but constructively about what they are doing, and theorists of organizations urgently need to figure out how to think in a similar fashion about how to cope with that fact.

Acknowledgments

The authors would like to acknowledge the support received from INSEAD's R&D department and the Reginald H. Jones Center of The Wharton School. They are also grateful for the many suggestions, comments, and critiques generously offered by Ron Adner, Max Boisot, Tim Folta, Anna Grandori, Bruce Kogut, Arie Lewin, Bart Nooteboom, Gabriel Szulanski, Christoph Zott, and three anonymous reviewers. Of course, the authors are responsible for any remaining errors and omissions. In its working paper version, this article was titled "From Organizational Routines to Dynamic Capabilities."

Endnotes

¹Of course, this is not to say that firms do not articulate and codify knowledge obtained from the external environment. As we will see in §2, environmental scanning activities are an important antecedent to generative variation processes. However, we argue that such activities are usually better understood as stimuli to the initiation of proposals to modify existing routines, rather than as mechanisms directly shaping the development of dynamic capabilities. To illustrate, a sound understanding of what competitors do and customers desire represents a crucial element of any firm's competitive strategy, but in and of itself does

not make it any more capable of creating and modifying its own set of operating routines. Also, operating routines typically involve tacit knowledge; hence they are highly unlikely to be developed or shaped simply by the observation of competitors, suppliers, customers, or other external constituencies. Such knowledge has to be developed "in house" through a set of activities and cognitive processes focused on the organization's own routines.

²This model has significantly benefited from many discussions had during the "Knowledge and Organization" workshop held at Warwick University. We would like to acknowledge the contributions of Arie Lewin, Anna Grandori, and Bart Nooteboom, among many other participants. The model is also closely related to Max Boisot's notion of the social learning cycle (Boisot 1998), which influenced our thinking and spurred some of the insights that follow.

³The justification for the explicit consideration of replication processes as distinct from the retention ones becomes considerably weaker as the nature and the size of the firm considered changes. Small entrepreneurial start-ups, for example, might necessitate replication processes to a far lesser (and perhaps negligible) degree, compared to a multinational, multidivisional organization.

⁴For those interested in this specific empirical context, please refer to an article entitled "Two grandmasters at the extremes," which appeared in the *Alliance Analyst* on Nov. 25, 1995.

⁵Kale (1999) is the only empirical work, to our knowledge, that has measured all the three mechanisms theorized about in this paper. Knowledge articulation processes are positively and significantly correlated with alliance performance, with an overall explanatory power comparable to the knowledge codification proxy, and larger than experience accumulation.

⁶See especially the discussion and references in Adler and Borys (1996) and empirical work cited in §4.

References

- Adler, P. S., B. Borys. 1996. Two types of bureaucracy: Enabling and coercive. *Admin. Sci. Quart.* **41** 61–89.
- Argote, L. 1999. *Organizational Learning: Creating, Retaining and Transferring Knowledge*. Kluwer Academic Publishers, Boston, MA.
- Argyris, C., D. Schon. 1978. *Organizational Learning*. Addison-Wesley, Reading, MA.
- Boisot, M. H. 1998. *Knowledge Assets: Securing Competitive Advantage in the Information Economy*. Oxford University Press, Oxford, U.K.
- Brown, J. S., P. Duguid. 1991. Organizational learning and communities of practice: Toward a unified view of working, learning and innovation. *Organ. Sci.* **2** 40–57.
- Brown, S. L., K. Eisenhardt. 1997. The art of continuous change: Linking complexity theory and time-paced evolution in relentlessly shifting organizations. *Admin. Sci. Quart.* **42** 1–34.
- Campbell, D. 1969. Variation and selective retention in socio-cultural evolution. *General Systems* **16** 69–85.
- Cangelosi, V. E., W. R. Dill. 1965. Organizational learning: Observations toward a theory. *Admin. Sci. Quart.* **10** 175–203.
- Cohen, M., P. Bacdayan. 1994. Organizational routines are stored as procedural memory: Evidence from a laboratory study. *Organ. Sci.* **5** 554–568.
- Cohen, W. M., D. A. Levinthal. 1990. Absorptive capacity: A new

- perspective on learning and innovation. *Admin. Sci. Quart.* **35** 128–152.
- Collis, D. 1994. Research note: How valuable are organizational capabilities? *Strategic Management J.* **15** 143–152.
- Cormier, S., J. Hagman. 1987. *Transfer of Learning: Contemporary Research and Applications*. Academic Press, San Diego, CA.
- Cowan, R., D. Foray. 1997. The economics of codification and the diffusion of knowledge. *Indust. Corporate Change* **6** 595–622.
- , P. David, D. Foray. 2000. The explicit economics of knowledge codification and tacitness. *Indust. Corporate Change* **9** 211–253.
- Dosi, G., R. Nelson, S. Winter, eds. 2000. *The Nature and Dynamics of Organizational Capabilities*. Oxford University Press, New York.
- Duncan R., A. Weiss. 1979. Organizational learning: Implications for organizational design. *Res. Organ. Behavior* **1** 75–123.
- Eisenhardt, K., J. Martin. 2000. Dynamic capabilities: What are they? *Strategic Management J.* **21** 1105–1121.
- Gavetti, G., D. Levinthal. 2000. Looking forward and look backward: Cognitive and experiential search. *Admin. Sci. Quart.* **45** 113–137.
- Gick, M., K. Holyoak. 1987. The cognitive basis of knowledge transfer. S. M. Cormier and J. Hagman, eds. *Transfer of Learning: Contemporary Research and Applications*. Academic Press, New York.
- Glynn, M. A., T. K. Lant, J. M. Miliken. 1994. Mapping learning process in organizations: A multi-level framework linking learning and organizing. *Advances in Managerial Cognition and Organizational Information Processing*, vol. 5. JAI Press, Greenwich, CT, 43–83.
- Haleblian, J., S. Finkelstein. 1999. The influence of organization acquisition experience on acquisition performance: A behavioral learning theory perspective. *Admin. Sci. Quart.* **44** 29–56.
- Henderson, R. M., K. B. Clark. 1990. Architectural innovation: The reconfiguration of existing product technologies and the failure of established firms. *Admin. Sci. Quart.* **35** 9–30.
- Holland, J., K. Holyoak, R. Nisbett, P. Thagard. 1986. *Induction: Processes of Inference, Learning, and Discovery*. MIT Press, Cambridge, MA.
- Holyoak, K. J., P. Thagard. 1995. *Mental Leaps*. MIT Press, Cambridge, MA.
- Kale, P. 1999. Building an alliance capability: A knowledge-based approach. Unpublished dissertation, University of Pennsylvania, Philadelphia, PA.
- Kogut, B., U. Zander. 1992. Knowledge of the firm, combinative capabilities, and the replication of technology. *Organ. Sci.* **3** 383–397.
- Leonard-Barton, D. 1992. Core capabilities and core rigidities: A paradox in managing new product development. *Strategic Management J.* **13** 111–125.
- Levinthal, D., J. G. March. 1988. A model of adaptive organizational search. *Decisions and Organizations*. Basil Blackwell, New York, 187–218 (originally published 1981).
- Levitt, B., J. G. March. 1988. Organizational learning. *Ann. Rev. Soc.* **14** 319–340.
- Lippman, S. A., R. P. Rummelt. 1982. Uncertain imitability: An analysis of interfirm differences in efficiency under competition. *Bell J. Econom.* **13** 418–438.
- March, J. G. 1991. Exploration and exploitation in organizational learning. *Organ. Sci.* **2** 71–87.
- , D. Levinthal. 1993. The myopia of learning. *Strategic Management J.* **14** 95–112.
- , H. A. Simon. 1958. *Organizations*. Wiley, New York.
- , L. S. Sproull, M. Tamuz. 1991. Learning from samples of one or fewer. *Organ. Sci.* **2** 1–13.
- Mukherjee, A., M. Lapré, L. Van Wassenhove. 1998. Knowledge driven quality improvement. *Management Sci.* **44** 535–549.
- Narduzzo, A., E. Rocco, M. Warglien. 2000. Talking about routines in the field: The emergence of organizational capabilities in a new cellular phone network company. G. Dosi, R. Nelson, and S. Winter, eds. *The Nature and Dynamics of Organizational Capabilities*. Oxford University Press, Oxford, U.K. 27–50.
- Nelson, R., S. Winter. 1982. *An Evolutionary Theory of Economic Change*. Harvard University Press, Cambridge, MA.
- Nonaka, I. 1994. A dynamic theory of knowledge creation. *Organ. Sci.* **5** 14–37.
- , H. Takeuchi. 1995. *The Knowledge-Creating Company*. Oxford University Press, New York.
- Ocasio, W. 1997. Towards an attention-based view of the firm. *Strategic Management J.* **18** 187–206.
- Selznick, P. 1957. *Leadership in Administration: A Sociological Interpretation*. Row, Peterson, and Co., Evanston, IL.
- Teece, D. J., G. Pisano, A. Shuen. 1997. Dynamic capabilities and strategic management. *Strategic Management J.* **18** 509–533.
- Tyre, M. J., W. J. Orlikowski. 1994. Windows of opportunity: Temporal patterns of technological adaptation in organizations. *Organ. Sci.* **5** 98–118.
- Walston, S. 1998. Does re-engineering really work? Extracted from an unpublished dissertation, University of Pennsylvania. Presented at the *EGOS Conf. Maastricht*, The Netherlands, July 1999.
- Weick, K. 1979. *The Social Psychology of Organizing*, 2nd ed. Addison-Wesley, Reading, MA.
- . 1995. *Sensemaking in Organizations*. Sage Publications, Thousand Oaks, CA.
- Winter, S. 1987. Knowledge and competence as strategic assets. D. J. Teece, ed. *The Competitive Challenge: Strategies for Industrial Innovation and Renewal*. Ballinger, Cambridge, MA, 159–184.
- , G. Szulanski. 2000. Replication as strategy. *Organ. Sci.* **12** 730–743.
- Zander, U., B. Kogut. 1995. Knowledge and the speed of transfer and imitation of organizational capabilities: An empirical test. *Organ. Sci.* **6** 76–92.
- Zollo, M., 1998. Knowledge codification, process routinization and the development of organizational capabilities: Post-acquisition integration in the U.S. banking industry. Unpublished doctoral dissertation, University of Pennsylvania, Philadelphia, PA.
- , J. Reuer. 2000. Experience spillovers across corporated development activities. INSEAD Working paper, Fontainebleau, France.

Accepted by Arie Lewin.