

Making Sense of Limited Learning¹

CONTROVERSIAL ISSUES OF ORGANIZATIONAL LEARNING

From our review of the organizational learning literatures, we draw the following issues which seem to us crystallize the most important current debates in the field.

The first issue centers on the problem of *levels of aggregation*. At what levels of aggregation – individual, interpersonal, group, intergroup, or whole-organization – does it make sense to speak of productive organizational learning? More specifically, how should we think about the role of interpersonal interaction and its relationship to patterns of activity at higher levels of aggregation? Some of the most prominent researchers in the scholarly branch of the field of organizational learning refer mainly to the actions or interactions of departments or groups within the world of an organization or of whole organizations within larger ecologies. In contrast, we emphasize the importance of interpersonal inquiry carried out within the constraining or enabling context of an organizational learning system, focusing on how such inquiry interacts with processes described as occurring at higher levels of aggregation. We adopt this emphasis in order to provide both knowledge useful to practitioners and robust explanations of the phenomena of organizational learning and nonlearning.

A second issue relates to the *meaning of “productive learning.”* Our treatment of this issue centers on the distinction between single- and double-loop learning, a distinction that has been rather widely adopted, in one form or another, by authors in both branches of the organizational learning literature. Some authors, especially economists and proponents of the learning organization, focus on the importance of instrumental, single-loop learning. Other authors, mainly among the scholarly skeptics, question the values served by initiatives taken in the name of the learning organization. Still other authors speak of the importance of double-loop learning but without serious attention to the behavioral conditions for its achievement. In addressing the issue of the meaning of productive organizational learning, we underline the importance of distinguishing between two kinds of organizational double-loop learning. The first kind involves change in organizational outcomes, especially in the values contained in an organization’s instrumental theory of action. The second involves change in the values of an organization’s theory-in-use for the process of organizational inquiry. How are these two kinds of double-loop learning related? Is it possible for an organization to achieve the first without achieving the second?

A third issue has to do with the nature of the *impediments to productive organizational learning that arise in real-world organizations*. Learning skeptics have discussed cognitive

limitations, the pull of past success, and the subversive effects of organizational politics. What is the likelihood that, in the face of such impediments, real-world organizations can achieve productive organizational learning, especially of the double-loop variety?

A fourth and final issue has to do with *the kinds of interventions that are likely to be effective* in enhancing organizational capability for productive learning. What can we expect from interventions based on the introduction of various sorts of organizational enablers: formal role and authority-structures (e.g., decentralization), information systems (e.g., activity-based costing), systems of incentives (e.g., high-performance organizations), procedures and systems for organizational inquiry (e.g., TQM and reengineering)? What can we expect from attempts to introduce new cultures, with their associated beliefs and assumptions? Each of these types of enablers has been advocated, at one time or another, by some proponents of the learning organization.

The four issues are distinct, but as they arise in particular cases, they are logically and empirically interconnected. In this chapter, we consider them through the lens of two groups of studies. We begin with two studies of naturally occurring innovation or adaptation: Van de Ven and Polley's (1992) study of technological innovation in the medical field, and Burgelman's (1994) study of the processes that led to Intel's "strategic exit" from the DRAM business. In these cases, we focus mainly on interactions between interpersonal inquiry within a limited learning system and processes of learning or nonlearning at higher levels of aggregation. We turn next to studies of attempts to promote productive organizational learning through the introduction of organizational enablers: a review of Cordiner's decentralization of the General Electric Company; a study by Argyris and Kaplan of the introduction of Activity-Based Costing; two retrospective analyses, by Schneiderman and Mallinger, of programs of Total Quality Management (TQM); George Roth's study of two cases of Business Process Reengineering; and Michael Beer's analysis of his Strategic Human Resource Management (SHRM) intervention.

RESEARCH STUDIES OF ORGANIZATIONAL ADAPTATION AND LEARNING

Van de Ven and Polley: "Learning While Innovating"

In a noteworthy study of technological innovation in the medical field, Van de Ven and Polley (1992) set out to test the familiar adaptive model of trial-and-error learning. In its simplest form, this model says that we persist in a prior course of action if we perceive its results as positive and deviate from that course if we perceive its results as negative. The authors propose to apply this model which had previously been explored in "unrealistic laboratory or simulation studies" to the "highly uncertain organizational field setting" of entrepreneurs in an innovation unit. (p. 113) In certain respects, they attempt to make a rigorous test of the model; but in other respects, at least as important, we find their treatment of organizational innovation and learning to be limited and unrigorous.

The authors tell two complementary stories of technological innovation. The first, articulated at a relatively high level of aggregation and couched in the language of statistical analysis, tests hypotheses derived from the adaptive learning model against the behavior of two organizational units, the "innovation's internal management team" and the "external resource controllers." The second story, presented in a running narrative (but not illustrated by actual dialogue), tells how individuals behaved at key points in the innovation's trajectory. But while the authors suggest some possible causes of the behavior that resulted in limited organizational learning, they offer no coherent theory to account for it, nor do their concluding prescriptions offer much of a basis for improving it. Their article demon-

strates the limits of both a high-level analysis of organizational learning and a reliance on quantitative, statistical methodology. One has the impression that, although the authors intuit the importance of their behavioral story, they have no way of explaining or getting underneath it.

Testing the adaptive learning model The subject of Van de Ven and Polley's (1992) longitudinal study is

a joint venture by three corporations to create a business by developing a new medical technology called therapeutic apheresis . . . which treats disease by removal of pathogenic blood components. (p. 98)

Apheresis had been around for some time, but it had not gained widespread usage because of its limited ability to remove specific components of blood. The Therapeutic Apheresis Program (TAP) centered on a device consisting of filters, pumps, and computer controls that was designed to make more specific separations possible. The joint venture partners sought to combine their respective competencies in separation techniques, pump technology, and marketing. They planned a three-phase development process:

- 1 a product to compete with existing apheresis technology;
- 2 an advanced product to treat specific diseases using advanced filtration modules;
- 3 the development of future apheresis technologies.

Van de Ven and Polley conducted a "real-time field study of TAP's development from October 1983 to July 1988" (p. 99). They divided the innovation's eight-year trajectory into three stages:

- 1 an initial three-year gestation period, culminating in a formal decision to initiate and fund the TAP innovation as a joint interorganizational venture;
- 2 a three-year expansion period aimed at producing a commercial TAP device, expanding the development program and meeting FDA requirements; and
- 3 a two-year contraction period when the TAP device entered the market, experienced product difficulties, and was terminated.

As the authors test their model of adaptive learning against the data of the TAP story, their findings vary with the three phases of development. In the gestation period they find actions and outcomes unrelated; in the expansion period, they find them negatively related; and in the contraction period, positively related.²

In the expansion period, the authors find "little or no learning by trial and error." Rather, negative outcomes "led directly to continuing with the prior course of actions" which had "no effects on subsequent assessments of positive or negative outcomes nor on changes in goals or criteria" (p. 104). The authors interpret this "direct negative relationship" between outcomes and subsequent actions as suggesting that "TAP's early development may have largely consisted of escalating commitments to failing courses of action (as observed by Ross and Staw, 1986)."

In the contraction period, the authors do see evidence of adaptive, trial-and-error learning. They observe that entrepreneurs showed a propensity "to select the course of action that was rewarded," whereas negative outcomes "triggered interventions by resource controllers, which resulted in changes in the innovation unit's course of action" (p. 104). These changes led, in turn, to shifts in outcome criteria, suggesting to the authors that "outcome criteria may have shifted largely to justify changes" in the course of action.

The authors' findings lead them to pose three puzzling questions (p. 106):

- 1 Why did trial-and-error learning not occur during TAP's expansion period?
- 2 What explains the dramatic shift from little or no learning during the expansion period to trial-and-error learning during TAP's contraction period?
- 3 When learning occurred, it appeared too late. Why was TAP's development terminated?

The behavioral story It is in their exploration of these puzzles that the authors turn to the phenomena they call "behavioral." They observe that during the expansion phase the innovation entrepreneurs were held accountable for achieving the overly optimistic plans they had presented in order to secure funding. This, in turn, triggered "'sugar-coated' administrative reviews" (p. 106). None of managers of the innovation team interviewed by the authors showed a willingness to "document uncertainties or to propose a more extended timetable for start-up, because they feared that would decrease their chances of obtaining startup funding." The top managers of the parent companies that funded TAP accepted the overoptimistic planning targets presented by the entrepreneurs, although they admitted privately that "they discounted certain projections as 'fluff' and expected the plan to change." Nevertheless, they held publicly to the conviction that TAP managers should be kept accountable to their plans. This, in turn, "precipitated [our emphasis] *the onset of a vicious cycle of impression management* between innovation managers and resource controllers" (p. 106).

While development efforts were largely successful in this period, they did encounter critical difficulties – for example, manufacturing defects and problems in scaling up production of the filtration module – which resulted in slipped schedules for market introduction, deferred sales revenues, and delayed development of the Phase II device (p. 106). During the administrative review sessions that were held every six months or so, the TAP managers reported information about these problems to the resource controllers, but they did so in ways that were calculated to discourage problem solving and learning. Van de Ven and Polley (1992) report that before each review session the TAP managers spent a day

rehearsing their presentations, developing tactics and scripts on how they would respond to possible questions of top managers, . . . preparing "slick" visuals . . . and [reconstructing] negative information in a positive frame, with assurances that they were in control of problems and presented action plans for addressing [them] . . . The resource controllers . . . relied on this indirect information to assess TAP's progress . . . and tended to mimic the TAP success criteria and concerns they were told by innovation managers. (p. 107)

In the expansion phase, setbacks occurred frequently, and red flags were raised with increasing frequency concerning TAP's development, schedule, directions, or financing. Nevertheless, the authors observed relatively few attempts to "detect these warnings as errors and correct the detected errors." Some warnings provoked discussion and debate, but most "were simply aired without response or were dismissed as irrelevant." At no time in the expansion period did TAP managers "seriously question the assumptions or validity of the technological paths they were pursuing" (p. 108).

The authors account for these remarkable (though classic) findings by reference to two main factors, in addition to vicious cycles of impression management. First, they state that a high level of discontinuity in the TAP personnel, due to normal job mobility and promotion in the parent companies, disrupted the continuity of attention necessary for adaptive learning. Second, they observe that the proliferation of "parallel and independent streams

of research activities" appeared to have the effect of "[masking] attention to the core innovation idea: the viability of producing a single reliable filtration device" (p. 107). They suggest that this noise, which went far beyond the information processing capacity of individuals, makes it easy to understand "why so few messages were detected as errors," and "why few attempts were observed to correct detected errors" (p. 110).

In the contraction phase, things changed. First of all, TAP got into serious trouble when its managers tried to introduce the Phase I devices into the European market and use them to conduct clinical trials of the Phase II filtration module. Even so, the authors report, the negative assessment of TAP's performance did not begin to grow until it was revealed at an administrative review session that development targets for the Phase II device had slipped by over a year and that sales revenues from the Phase I device were at 10 percent of projections. Only then did resource controllers begin to raise serious questions about TAP's development and intervene to break the pattern of persevering in failing courses of action. At this point the resource controllers put new TAP managers in charge of manufacturing. These managers suspended all shipments of Phase I devices and conducted an extensive internal review and audit of the scaleup and manufacturing problems; as a result several TAP people were laid off. As one member of the TAP unit observed at the time, "The honeymoon is over!" (p. 112).

These events produced what the authors regard as a further cause of increased adaptive learning, that is, a concentration of attention on pressing operational problems which had the effect of "reducing complexity and focusing attention on a single issue," which, in turn, "facilitated trial-and-error learning" (p. 110). Sophisticated technological and engineering trials and adaptations were made, and long-standing manufacturing defects were finally corrected. As a result, the TAP device was found to achieve "a reliability rate comparable to competing apheresis devices in the market" (p. 112).

Nevertheless, the troubles experienced in this phase shook the resource controllers' confidence. They began to pay serious attention to the mounting development costs that were draining resources from alternative investments. In contrast, the innovation managers saw the recent setbacks as temporary and maintained their long-term commitment to the program (p. 112). The two parties developed "alternative stories that interpreted the same experience quite differently." Meanwhile, in the external environment, "latent forces" were gathering. In the parent companies, budgets were being reduced in the face of negative corporate earnings. A CEO, who had been one of TAP's principal supporters, died unexpectedly. There was news of technological advances by TAP's Japanese competitors and rumors that a new drug might do away with the need for apheresis hardware. Then, when it became evident that development targets had slipped by over a year and that market trials had revealed serious manufacturing defects, these latent forces were unleashed. In November of 1988, TAP was terminated.

In actuality, the authors point out, there were fewer threatening events during the contraction period than in the expansion period, but TAP had lost its credibility. The resource controllers launched a search for a new corporate investor, and, finding none, they declared TAP a failure. Was it a failure? The authors end their story by quoting a senior executive in one of TAP's parent organizations: had the program been allowed to continue, he thought, it might have succeeded.

The TAP analysis analyzed As an account of a technological innovation, the TAP story has a familiar ring (see, for example, Schön, 1967). Throughout the early stages of the process, researchers and general managers pursue a game of reciprocal deception and attempted unilateral control: researchers smooth over negative information, and general managers delay in attending to negative signals. This goes on until a crisis occurs, at which

point the researchers can no longer mask the negative information, managers can no longer avoid paying attention to it, and they overreact.

At the very end of their version of this classic story, the authors suggest some things organizations can do to increase the odds of adaptive learning.

- 1 separate the planning from the funding of a new venture, so as to avoid optimistic estimates aimed mainly at securing funding (p. 114);
- 2 find ways to decrease the bull index, i.e., the high amount of impression management that goes on between innovation teams and investors...and distorts information needed for trial-and-error learning;
- 3 keep the core innovation team intact during innovation development;
- 4 keep innovation development and business creation separate, so as to reduce information overload (p. 114); and
- 5 "alter the agendas and formats of...administrative review sessions in order to increase the amount of candid information exchange and learning" (p. 115).

What is striking about Van de Ven and Polley's (1992) treatment of the TAP story is that they are so rigorous in their testing of the adaptive learning model, yet they make so little explanatory use of their behavioral data. The mixed results of their test of the adaptive learning model lead to important puzzles.

- Why was adaptive learning so long delayed?
- Why was there a dramatic shift to adaptive learning in the contraction phase?
- Why, when adaptive learning finally occurred, did controllers pull the plug?

In their use of the behavioral story to explain these puzzles, the authors note the ambiguous information, the shifting objectives, the uncertainties surrounding the development, complexity and noise in the information environment, and the discontinuities of personnel. They treat these explanations as partial causes of the lack of adaptive learning in the expansion phase.

But such phenomena are the givens of research and development programs. All such programs encounter them in greater or lesser degree. The key question is how and through what kinds of interpersonal inquiry participants in a development deal with such phenomena. Given the uncertainties inherent in a development project, do the participants seek to make public their doubts as well as their beliefs? Do they invite challenges to their positions or surface the dilemmas with which they themselves are struggling? Given endemic noise, ambiguity, and uncertainty, do they seek out and invite public testing of assumptions that underlie the core innovation idea? Or do they enter into a vicious cycle of impression management, as in the TAP story, with innovation managers suppressing uncertainties and smoothing over negative information in order to keep the controllers from taking setbacks seriously and with controllers colluding with the researchers by holding them publicly to overoptimistic targets whose fluff they privately acknowledge?

The authors note these features of organizational inquiry, which they express at relatively high levels on the ladder of inference ("sugar coating," "impression management"), but they offer no coherent theory to explain them. From our point of view, the participants' vicious cycles of impression management are predictable variants of primary inhibitory loops, triggered by the threat or embarrassment associated with situations of uncertainty and ambiguity. In these collusive processes, both parties employ a Model I strategy of mystery and mastery. They keep mysterious what they know to be true (the fluffy estimates, the likelihood that plans will change), while outwardly they try to master the responses of

the other party. Under these conditions, errors and setbacks in the development's trajectory and uncertainties about its future cannot be made public; to do so would reveal both parties' deceptions. So the crucially important errors and uncertainties along with their undiscussability are kept undiscussable, until a critical event (the flawed entry into the European market, for example) makes error unavoidably visible, provoking a crisis that triggers the resource managers' precipitous termination of the project.

In our view, such a vicious cycle is explainable, indeed, predictable, on the basis of Model I theories-in-use, which shape interpersonal inquiry under conditions of threat or embarrassment. These theories-in-use reinforce and are, in turn, reinforced by O-I learning systems. As we see it, this is the systematic explanation of the patterns of unreliable organizational inquiry through which the actors in the TAP story respond to uncertainty, complexity, ambiguity, and discontinuity – the givens of technological entrepreneurship. From this viewpoint, one sees the flimsiness of the authors' prescriptions for enhancing adaptive learning. How will the separation of planning from funding decisions encourage adaptive learning if, once the funds are granted, the innovation team persists in its strategies of control through deception, and the resource controllers collude in appearing to swallow such deceptions in order to control the researchers by holding them publicly accountable for their formal targets? How will the participants reduce the 'bull index' when their theories-in-use lead them to produce and accept the bull? How will altering agendas and formats lead to an increase in candid information exchange and learning when such candidness would threaten to surface data that both parties actively work at keeping undiscussable?

The patterns of behavior implicit in the vicious cycle of impression management are unlikely to change unless the participants invite an intervention that creates the conditions for the kind of 'open, fact-based conversation' that Van de Ven and Polley recommend, or unless the participants learn to develop Model II theories-in-use, as we have illustrated earlier, so as to create such conditions for themselves. Model II action strategies – for example, the public testing of privately held assumptions and attributions – do not require a prior removal of conditions of ambiguity and uncertainty. On the contrary, it is precisely under profound uncertainty and ambiguity, when one cannot know what truth is, that Model II inquiry is most needed and most likely to create the conditions for good dialectic.

Whether or not readers accept our approach to the analysis of the behavioral phenomena described in the TAP story, we hope it is now clear why such phenomena must be considered characteristically organizational and crucially linked to organizational learning. Consider the following propositions arrayed in descending order on a ladder of aggregation:

- 1 In the expansion phase of the TAP project, there was an organizational failure of adaptive learning.
- 2 In their administrative reviews during this phase, innovation managers and resource controllers engaged in a vicious cycle of impression management.
- 3 In any given instance of such a review, an innovation manager, A, gave a 'sugar-coated' account of progress, knowing that he was doing so, but keeping that knowledge to himself; and a resource controller, B, privately discounted that statement as fluff, but kept that knowledge to himself.

If, following Van de Ven and Polley (1992), we treat item 1 as a proposition about organizational phenomena and item 2 as a partial explanation of item 1, then how can we not also treat item 2 as about organizational phenomena? And if we treat item 3 as a specification

and partial explanation of item 2, then how are we justified in excluding it from the set of propositions that refer to organizational phenomena? We cannot validly exclude from the class of propositions about organizations those that offer explanations of phenomena that we already treat as organizational. More specifically, we cannot validly exclude from the class of propositions about organizational learning those propositions about interpersonal inquiry in organizations that are adduced to explain observed patterns of organizational learning.

If this argument holds for the explanation of organizational learning, it holds all the more so for prescriptions aimed at enhancing it. Practitioners will be unable to act on such prescriptions unless they already have the skills required to carry out the actions prescribed. They will not be able to enact a prescription such as "Cut the bull index," or "Hold candid, fact-based conversations," unless the prescription is formulated in terms that enable them to envisage actions required for its implementation.

Burgelman: "Fading Memories"

In this article, Robert Burgelman (1994) reports on his study of Intel's responses, roughly between 1971 and 1991, to the changing competitive environment of its dynamic random access memory (DRAM) business and details the processes that led to Intel's eventual exit from that business. Like Van de Ven and Polley (1992), Burgelman tells a story of mixed success and failure in organizational learning and stresses the "inertial" lag of formal corporate strategy in relation to shifting conditions in the competitive environment. He differs from Van de Ven and Polley, however, in his conceptual framework, which is not adaptive learning but a theory about the role of strategy in firm evolution. Burgelman sets out to provide

insight into how the internal selection environment mediates the co-evolution of industry level sources of competitive advantage and firm-level sources of distinctive competence and into the link between corporate strategy and strategic action. (p. 24)

Burgelman tells the story of a 20-year transformation in Intel's strategic practice, tracking its passage through three stages:

- 1 an equilibrium state, harmony of dominant beliefs and strategic context, in which Intel thought of itself as a "memory company" (and more specifically, a DRAM company);
- 2 a transitional period of disharmony in which the basis for competitive advantage in the memory business changed, while Intel's formal strategy remained inert and momentum developed internally toward the opening up of new business opportunities; and
- 3 a later harmonious state in which Intel came to think of itself as a microprocessor company.

Ultimately, as Burgelman reveals, Intel learns. It learns to recognize that, through its incremental responses to challenges, mainly from Japanese competitors, it has dealt itself out of the DRAM business and should not try by massive capital investment to reestablish its position in that business. It learns to exploit the technological opportunities that came about as byproducts of its experience in the DRAM business (the EPROMS development pioneered by Dov Frohman, which came through exploring reliability problems with Intel's MOS technology; and the new chip architecture, which grew out of a customer's request for

a new chip set). These opportunities, initially recognized or tolerated by top management, were, according to Burgelman, internal variations. They were eventually selected by Intel's internal environment, and they eventually contributed to the technological basis for its microprocessor business. Intel also learns to shift its scarce manufacturing capacity from DRAM to microprocessor technology, and Intel's top management eventually learns to redefine Intel as a microprocessor company.

But, as Burgelman tells it, the Intel story is also one of failure to learn or of inertial delays in learning. Burgelman frames this side of the story in terms of six puzzling questions around which he structures his analysis. We focus on four of these.

- 1 Why did Intel, the first successful mover in DRAMS, fail to capitalize on and defend its early lead?
- 2 How did it happen that the bulk of Intel's business had shifted away from DRAMS, and DRAM market share was allowed to dwindle, while top management, even in 1984, was still thinking of DRAMS as a strategic business for the company?
- 3 How was it possible that middle-level managers could take actions that were not in line with the official corporate strategy?
- 4 Why did it take Intel's top management almost a year to complete the exit from DRAMS after the November 1984 decision not to market 1 Meg DRAMS? (pp. 29–30)

These questions point to a series of strategic mistakes that Intel seemingly made and to inertial delays in its attempts to reconcile its formal strategy to changing competitive conditions and to convert its new strategic intent into action. Such mistakes and delays are to be understood as failures in the timely detection and correction of significant errors (surprises or mismatch of outcomes to expectations). However, there is some ambiguity and perhaps ambivalence in Burgelman's treatment of his puzzles. He certainly does frame them in terms of Intel's apparent failure to take effective, timely action or to match its thought to a changing reality. But he also suggests that in certain respects Intel's delayed responses were functional. He speculates, for example, that Intel

would probably have done worse if it had simply divested the DRAM business and entered the new business through acquisition . . . [because it would have thereby] failed to capitalize on the full potential of its distinctive competencies in DRAMS, [some of which] could be effectively deployed in the microprocessor business . . . (p. 52)

This ambiguity plays an important role in the discussion that follows.

Burgelman's approach to his puzzles has both an historical and a systems dimension. He argues that Intel secured a strong competitive advantage in the early days of the DRAM business on the basis of its distinctive technical competence in designing, building, and manufacturing the DRAM technology. He argues further that the early successes of DRAM caused Intel's corporate strategy to be dominated by its DRAM technology. He postulates a microtheory of success and failure along the lines of March's "competence trap": When a pattern of factors is clearly related to a firm's understanding of its success, this pattern tends to be reinforced and thus to persevere, even after it ceases to be effective in the competitive business environment.

Burgelman describes several phenomena that reinforced established patterns associated with Intel's earlier success:

- the self-interest of business unit heads who had considerable latitude for decision-making yet had no interest in putting themselves out of business;

- the “bounded rationality” that kept top or middle managers from anticipating the competitive market dynamics of product and manufacturing technology that would squeeze out Intel’s share of the DRAM market; and
- top management’s “emotional attachment” to the DRAM business, which made it reluctant to get out of that business, just as Ford would be reluctant to “decide that it should get out of the car business.” (p. 41)

In contrast to the factors that kept Intel’s top management from letting go of its attachment to the DRAM business and its image of Intel as a “memory company,” Burgelman postulates an “internal selective environment” made up of “structural and strategic contexts shaping strategic actions” on the part of middle managers. In the Intel case, middle managers used their discretionary freedoms to take actions that had the effect of opening up new business opportunities. For example, they pursued the development of technological offshoots of the DRAM business that would lay the groundwork for microprocessors. And they allocated scarce manufacturing capacity to favor microprocessors over DRAMs, making use of the established corporate rule that manufacturing capacity should be allocated so as to “maximize margin-per-wafer-start” (p. 43), even while top management continued to support the view that Intel should remain in the DRAM business. In these and related ways, middle managers gradually undermined the position of the existing DRAM business and helped “dissolve,” as Burgelman puts it, the earlier strategic context of the firm.

Burgelman calls his perspective “inside-out.” He tells how Intel’s external business environment created conditions – changes in the market and DRAMs becoming a commodity – to which Intel responded with decreasing effectiveness. But he treats the “inside” as crucial. Unlike other writers, especially economists, who relegate inside activities largely to a black box, Burgelman opens up the box. He sees that as Intel sought to maintain its competitive advantage by acting in ways that were consistent with its perception of its distinctive competence, it did and failed to do things that actually led to its competitive disadvantage. For example, although Intel’s top managers did realize that DRAM had gone from a premium-priced to a commodity product (the price signals were rather clear), they persisted in believing that they could regain the lead by applying their traditional strengths in process technology to come up with innovative products that would be premium-priced niche products at first but would have to be adopted eventually by the entire market. The moment DRAMs became a commodity, Intel’s strategy of being the first to introduce premium-priced products into the market became outmoded. Yet top management acted as if this were not the case.

In addition, Burgelman suggests that middle managers believed top management was closed to constructive confrontation about their emotional attachment to the DRAM business. From our perspective, the middle managers were making attributions about top management’s openness to learning. From the data in the case, we infer that they never tested these attributions. In a discussion with Argyris, Burgelman confirmed this inference. He said, in effect:

Given the status and power of the process development people, anyone who went into a meeting to propose a view different from theirs would be unlikely to carry the day. Furthermore, top managers were seen by middle managers as believing that Intel could not make changes away from DRAMs. Top management was seen as unsure.

This behavior provides a partial explanation of top management’s “blindness”: they were never confronted with data that would help them realize the impact of their behavior. Burgelman’s explanation focuses on the top management’s uncertainty and, in the face of that uncertainty, its resistance to change. We would focus, in addition, on middle managers’

actions that helped to keep top management from realizing the full implications of the shift in the DRAM business, which, in turn, let top managers act in ways middle managers explained as resistance to change. Burgelman focuses on the self-sealing processes of the top managers, not on those created by the middle managers. If an intervention were designed to interrupt the limited learning processes revealed by the Intel case, it would fail if it were one-sided.

Burgelman reports that Intel's top management encouraged constructive confrontation. Andy Grove (then COO, now CEO) told managers at all levels that decisions should be data-driven and that power and emotional biases should play no role in decision-making. We hypothesize, nevertheless, that middle managers' attributions to the top were not openly discussed and analyzed in ways that would test their validity and allow them to be corrected if they proved mistaken.

Burgelman describes an additional domain of undiscussables. He notes that when DRAMs became a commodity, manufacturing capability became the dominant success factor. Yet process development (TD) people continued to downgrade manufacturing and distance themselves from it. TD people continued to frame problems in mainly scientific terms, treating manufacturing people as "tweakers." The manufacturing people reacted by developing a "not invented here" (NIH) syndrome. Gordon Moore, then CEO, reacted to this interdepartmental rivalry by placing both groups in geographical proximity in order to foster greater cooperation between them. Moore had experienced similar problems at Fairchild and wanted to avoid them now at Intel, an attempt to learn from his earlier experience. Burgelman reports that this strategy was at first effective but became much less effective when competitive pressures made the differences between TD and manufacturing much more salient.

Burgelman also observes organizational defenses at work in the relationship between upper-level middle managers and the Board. By 1980, when the company's total market share in DRAMs was less than 3 percent, these managers saw the handwriting on the wall. Yet they would learn of board meetings where Moore, then CEO, and Gelbach, the head of marketing, would defend continued expenditures on DRAM - Moore, because he saw DRAMs as the company's "technology driver" and Gelbach, because he believed Intel had to offer its customers a one-stop semiconductor shopping list. Andy Grove (then COO) remained silent. As a result of such board-level interactions, middle managers were frustrated.

These multileveled undiscussables reinforced by organizational defensive routines would make Grove's advocacy of "constructive confrontation" sound more like espoused theory than theory-in-use, namely, his view that "one of the toughest challenges is to make people see that their 'self-evident truths' are no longer true," or "organizations should practice creative destruction of routines that are no longer effective."

Burgelman told Argyris that top management recognized the organizational defensiveness which they saw as existing primarily at the level of some middle managers. As far as we can tell, top management did not examine this defensiveness directly and forthrightly with the middle managers. If so, we have top managers acting to reinforce middle managers' unawareness of their impact on those at the top.

Top management attempted to deal with interdepartmental conflicts by bringing departments into geographic proximity, defining rules to reduce the conflicts, and providing incentive systems that would do so. Over the long haul, none of these strategies worked, although some of them were effective in the first instance.

Burgelman's analysis analyzed We would say that there is only partial truth in Burgelman's hypothesis that top management persisted in its commitment to the DRAM

business because of emotional attachment to it and uncertainty about its importance. What is needed is an explanation of how subordinates colluded with their superiors to create domains of undiscussables that would inhibit learning around these issues.

Like Burgelman, we also try to understand how it comes about that organizations do not detect and correct significant errors, but we seek to go further. We want to understand how it comes about that Intel or other organizations are unaware that they are unable to detect and correct significant errors. The inability to detect and correct error is skilled because we can connect it to a theory-in-use; hence, it is "skilled incompetence". The existence of skilled incompetence means the unawareness of the inability, and unawareness is also connected to a theory-in use – hence, it is "skilled unawareness".

We would ask, for example, what is included in this concept of "emotional attachment"? What is the theory-in-use that produces such behavior? Is Burgelman using the concept to mean that emotional attachment always causes blindness? If not, what would be the difference between the two types of emotional attachment?

We would add an additional puzzle to those queries with which Burgelman begins his paper: How could the conditions named in his four puzzles hold when top management espoused a theory of management intended to prevent them from occurring?

To solve this puzzle would require much more directly observable data about what was actually said when people at Intel discussed these issues. For example, how did the middle-level managers craft their conversations? Did they ease in so much that their points were obscured? Or did they become so confrontational that the others could discount their views because they were obviously emotional?

Did the middle managers engage top managers in a dialogue about the attributions they were making about them? Did they explore their attributions that these issues were undiscussable, and their undiscussability also undiscussable? Did they explore their attributions that even if these issues were discussed, top management would be uninfluenced? Did they ever explore whether the top managers were unaware of the "blindness" that middle managers attributed to them?

Burgelman states that Grove did not at first believe that the lag time was as long or the inertia as deep as Burgelman concludes they were. To his credit, however, Grove conducted his own investigation and found that Burgelman's estimate was better than his own. We might ask, then, what Grove actually did. He questioned the relevant players and, months after the decision to exit the DRAM business, they gave him information they had not told him or that he had not listened to before.

Burgelman's explanations of the puzzling phenomena of delayed or failed organizational learning do not account for such phenomena developing in an organization in which cultural context, top management support, and the sheer brightness of the Intel managers would seem to militate against their occurring.

If Intel failed to capitalize on and defend its early lead in the DRAM business, for example, what actually produced the failure? To say that it was top management's emotional attachment to DRAMs is to pinpoint a small group within an abstract explanation. How did top management actually behave? How did the middle-level managers confront top management's emotional attachment, if they actually did? How did they attempt to change it? In other words, we hypothesize that the puzzling phenomena of delay and inertia were caused both by top management's emotionality and blindness and by middle managers' ways of dealing with them, that is, by bypassing these issues and acting as if they were not doing so.

We would also suggest that an operational definition of a more complete explanation of the puzzles could be used to design actions to change the causes of inertia. The causes Burgelman identifies – such as top management's emotional attachment to DRAMs, interdepartmental rivalries between TD and manufacturing, self-reinforcing processes

that maintained outmoded strategic frames – are stated at a level of abstraction that cannot be used to change the organization or to enable the organization to learn to detect and correct its errors.

In our view, the fifth puzzle we introduced, how it came about that Intel's management was unaware of its inability to detect and correct its errors, underlies the other four.

Burgelman concludes his analysis by advancing a set of propositions, some of which, in the light of our own analysis, we find misleading. His first proposition states:

The stronger a firm's distinctive technological competence, the stronger the firm's tendency to continue to rely on it in the face of industry-level changes in the basis of competitive advantage. (p. 48)

We would add, however, that this proposition holds or holds more strongly whenever organizational defensive routines at top- or middle-management levels or between departments reinforce the failure to face up to industry-level changes and to legitimize the bypassing and cover up of such reinforcement.

Burgelman's third proposition states:

Firms whose internal selection criteria accurately reflect external selection pressures are more likely to strategically exit from some businesses than firms whose internal selection criteria do not accurately reflect external selection pressures. (p. 50)

We would add:

Firms whose internal selection criteria reflect and deal with external selection pressure should be observed to have minimal defensive routines (of the sort we have described earlier).

Burgelman's sixth proposition states:

Firms that have strategically exited from a business are likely to have a better understanding of the links between their distinctive competences and the basis of competition in the industries in which they remain active than firms that have not strategically exited from a business.

We believe this proposition is valid under the condition that organizational defensive routines dominate the examination of errors that are embarrassing and threatening to key players. It should not hold for firms that are not dominated by such defensive routines; these firms should have a better understanding before, rather than after, exit.

In sum, Burgelman's analysis of the Intel case focuses on puzzles rooted mainly in the organization's failures to detect and correct error:

- its failure to follow up its early lead in the DRAM business;
- its delay in matching its strategic ideas and self-image to the changing business reality; and
- its inertia in completing its exit from the DRAM business once it had decided to do so.

Burgelman's explanations of these puzzles hinge on the phenomena of top management's emotional attachment to the DRAM business, its bounded rationality and blindness concerning the mismatch of its strategic context to the changing business reality, the power and self-interest of business-unit managers associated with the DRAM business, and the rivalry between technological development and manufacturing departments. But each of these causes raises questions related to another puzzle: how did it happen that Intel's managers

were unaware of their inability to detect and correct their errors? This question points, in turn, to the layer of organizational phenomena that we regard as critically important both to the explanation of existing patterns of organizational learning and to the design of interventions aimed at changing those patterns, that is, the organizational defensive routines, and the O-I processes that constrained interpersonal inquiry between top- and middle-level managers, and between departments, thereby reinforcing a lack of awareness of the gap between Intel's strategic context and its changing business reality.

In a private communication to the authors in which he responds to an earlier version of our discussion of his article, Burgelman stresses the uncertainty with which Intel's top managers were dealing at the time they were grappling with the possibility of exiting from the DRAM business. He argues that their inertial delays and the defensive routines that reinforced those delays were actually functional since they enabled Intel to exploit in their microprocessors certain technological advances made by the DRAM process developers and also to prevent TD/DRAM people from leaving the company. He observes, in part:

I like to emphasize that one of the key points of my study is that it is very difficult for top and middle managers to examine at length what the strategic situation is (that is) faced by an organization in very dynamic environments. So much is going on simultaneously that the kind of exhaustive "airing out" of the strategic situation is probably unachievable. So, while I believe your hypothesis is a plausible one and that perhaps more effort should be spent on airing out strategic situations, I also believe (1) that we do not know enough yet about how "defensive routines" come about (perhaps my study contributes to precisely that!) and (2) that it is improbable that there is no cost at all associated with removing defensive routines. I submit that trying to remove defensive routines altogether might very well paralyze organizations operating in dynamic environments.

In this passage, Burgelman argues, as he further notes, that in dynamic and uncertain environments there may be considerable value in "strategic neglect."

The weakness we find in this argument is the following. We agree that inertial delays in a firm's response to a mismatch between its formal strategy and its actual strategic situation may be retrospectively discovered, on occasion, to be functional. But these delays may just as easily prove to be dysfunctional. How are managers to distinguish between functional and dysfunctional strategic neglect so long as they keep themselves unaware of their discordant beliefs and keep the crucial and threatening issues undiscussable?

Why, moreover, should we assume that the opening up of defensive routines would make it impossible to achieve deliberately the same benefits as those that were inadvertently achieved (through lack of awareness)? For example, Intel's top management could have chosen to keep its TD/DRAM capability on the very grounds of its uncertainty about its possible future utility, even in the face of a clear decision to exit the DRAM business.

Why then should we assume that the airing out of the strategic situation would have to be exhaustive (certainly, a recipe for paralysis!) when it could be limited to just those dilemmas that constituted the main bone of contention between top- and upper-level middle managers? One need not know just how defensive routines come about (although we believe our theory offers some insight into that question) in order to see how they constrain productive inquiry into critically important strategic issues.³

Main Points Drawn from Our Analysis of the First Two Studies

Reviewing our discussion of these two retrospective studies of "naturally occurring" organizational learning, we find:

- 1 Both studies show mixed results. There are some examples of productive organizational learning and some striking limits to productive organizational learning, both of which need to be explained.
- 2 Both studies illustrate “dynamic conservatism,” that is, an organization’s persistence in adhering to past patterns of practice in the face of information that should have caused them to change.
- 3 The researchers interpret their findings of limited productive organizational learning in two main ways:
 - the cognitive limitations due to “finite information processing capacity” experienced by organizational actors in the face of complexity, noise, ambiguity, and uncertainty; and
 - the intraorganizational games of interests and powers that evolve around an organization’s dominant action strategy and task system.
- 4 Our criticism of these explanations is that they only partly account for the first-order errors the researchers set out to explain and do not at all deal with the additional puzzle we find in second-order error: the members’ being unaware of their inability to detect and correct first-order error, or their limited ability to act on what they knew.
- 5 We reinterpret the researchers’ findings in terms of organizational defensive routines based on Model I theories-in-use and O-I learning systems which shape and constrain prevailing patterns of organizational inquiry.
- 6 Our analysis of the researchers’ studies addresses the “ladder of aggregation” problem: The dimension of interpersonal inquiry to which the researchers give only slight or ambivalent importance is critical to an adequate explanation of both the first-order errors the researchers identify in their cases and the second-order errors that we have identified in them.

STUDIES OF INTERVENTIONS AIMED AT PROMOTING PRODUCTIVE ORGANIZATIONAL LEARNING

We turn now to studies of a different sort: analyses of attempts to introduce *organizational enablers* of several different kinds with the aim of enhancing organizational capability for productive learning.

Revisiting Cordiner’s Decentralization of GE

The purpose of decentralization is to reduce the inefficiencies and rigidities caused by centralization. Alfred Sloan decided in the 1920s that General Motors was not likely to grow or to be managed effectively if it took, as he found out, nearly a year to get replies to many of his memos. Ralph Cordiner arrived at similar conclusions about GE in the 1950s when he learned during the antitrust proceedings on light bulbs that many of his senior executives were withholding information from him and covering up that they were doing so.

Sloan, Cordiner and many other executives throughout the world attempted to correct this situation by asking what kinds of power and reporting relationships and what kinds of incentive schemes would encourage executives to be more candid with each other and with their superiors. Sloan and Cordiner, as well as many other executives, looked for the answer in new organizational structures and in reward-and-penalty processes. The fundamental assumption they held in common, and one held by many prominent scholars of

organization, was that changes in structures and rules (structural “enablers”) can cause desired changes in behavior to occur.

Consistent with this assumption is the rule: place responsibility for action and authority to take action as close as possible to the players who have the relevant information. It is these individuals who should be given the opportunity to make informed decisions and choices. It is they who should be held responsible for implementing their choices and monitoring their effectiveness.

At GE in the 1950s, Cordiner and his colleagues espoused a theory of organizational decentralization based on rules that incorporated Model II values of valid information, informed choice, and personal responsibility for implementing choices and monitoring their effectiveness by detecting and correcting errors. We will demonstrate that GE was able to achieve limited double-loop learning outcomes, but their achievements were undermined over time as defensive routines intervened. The GE case clarifies the relationship between double-loop learning at the level of outcomes in organizational theory of action and at the level of processes of organizational inquiry, as it shows how persistence of O-I organizational inquiry undermined and limited double-loop organizational learning over time.

Cordiner realized that decentralization would require a change in culture and a change in the mental maps executives held about effective leadership. Accordingly he created one of the first executive development universities at Crotonville, New York. The first participants in the Crotonville program were individuals who were general managers of a business unit. Later, individuals who reported to these managers attended. The divisional heads to whom the general managers reported and the group heads to whom the divisional heads reported attended a short version of the month-long course. We will focus on the general managers and their superiors. However, we claim that the strength and difficulties that emerged in the relations between these two levels were replicated at all the other managerial levels.

Cordiner's senior vice president, Harold Schmidy, wrote a set of volumes that served as foundation for the Crotonville curriculum. The overwhelming emphasis of these volumes was to define the organizational features of decentralization. First there was an emphasis on defining correct reporting relationships. Then there was the emphasis on defining rewards and penalties to make the reporting relationship, in Cordiner's words, “come alive.” Cordiner and Schmidy, as well as the faculty, realized that “coming alive” meant that the managers' theories-in-use had to be changed. However, the fundamental assumption was that if the structure is defined correctly and correct rewards and penalty processes are in place, then the desired changes in managers' thought and action would occur.

What follows represents the recollections of one of the authors (Argyris) who helped to design the first program for general managers, taught in the programs, and was given the opportunity to observe some of the general managers when they returned to their organizations. Both Argyris and coauthor claim that what follows is not limited to GE but is true of most executive programs in which they and many of their academic colleagues have participated.

The theory of instruction at Crotonville was centralized in that the faculty was expected to teach and the general managers were expected to listen and to confront the faculty whenever they wished to do so. The material the general managers received represented the faculty's best efforts at defining the desired pattern of decentralization and the new controls that would make it happen. The fundamental intention was to specify decentralization and its many features as precisely as possible and to illustrate how these features should be implemented in the back-home situation. The faculty told the general managers that these plans were not likely to be correct or complete and invited vigilant on-line learning.

For the sake of illustration, let us begin with those general managers who learned the new concepts, policies, and skills. We refer to those who made the new espoused theory of action part of their theory-in-use for action within the organization. Argyris did find a few managers who hesitated, either because they did not believe top management truly believed in decentralization or because they did not feel competent and confident in the exercise of their new skills. But most general managers did make a serious effort to implement decentralization as it was taught to them. For example, if they could now sign checks for \$200,000, they did so. If they could decide whether to build or rent within some broad dollar limits, they did so. If they were free to hire new functional managers, they did so. Argyris observed many general managers who acted in ways that were consistent with the delegation of decision rights and exercised their newly given autonomy and responsibility.

These represent double-loop outcomes, although the managers were given no education that helped them to develop Model II skills. Note, however, the conditions for success. The general managers had the skills to sign checks, they had been educated in the new purchasing procedures, and they knew how to hire functional managers. In other words, where there was a relatively high degree of seamlessness between the concept of decentralization and the general managers' ability to produce it, the Model II governing values built into GE's theory of decentralization were translated from espoused theory to theory-in-use. In our opinion, this does represent double-loop change at the level of organizational theory-in-use without double-loop change at the level of interpersonal inquiry.

Our theory would predict that this success should persist as long as there were no major constraints or errors. As time passed by, however, limitations and errors did come to the surface. For example, no managerial policy for purchasing or for anything else could be so complete that it covered all contingencies. If the gaps between abstract policies and actual behavior did not require changes that were embarrassing or threatening, they were successfully bridged. If the gaps or their solutions proved embarrassing or threatening, they were handled by Model I action strategies. These strategies led to the general managers taking unilateral action, subordinates going under cover, and both groups acting as if this were not the case. For example, the subordinates argued that their newly given autonomy was insufficient, while their superiors often held the view that the subordinates were failing to take advantage of what freedoms they had. The embarrassment or threat involved in engaging these issues was bypassed, and the bypass was covered up.

Argyris observed general managers leading many meetings at their home sites. Whenever embarrassing or threatening issues surfaced, Model I theories-in-use, as well as predictable organizational defensive routines, were activated. General managers were observed dealing with the "difficult concerns" of their immediate reports with such comments as, "Look, let's give it a try" or "I know we can make this work if you folks would have a more positive attitude." The subordinates saw such actions as discrepancies between the organization's managerial theory of decentralization, which was consistent with Model II governing variables, and the actions of superiors, especially when managers were facing embarrassing or threatening situations.

The general managers were concerned about certain important cases in which they were getting mixed messages from their superiors. In a situation of ambiguity and tension about authority, the superior might say, "Of course you have the authority, but on this one it would be a good idea to check with me because policy questions are involved." Typically, faculty would respond to such examples by saying things such as, "Be positive and be patient" or "Show your superiors that the new rules do give you the authority to decide." This kind of advice bypassed both the mixed messages general managers were getting and the underlying ambiguities that frustrated them.

The consequence frequently observed was that such advice led to the violation of the concept of decentralization precisely when it was most needed, namely, when there were ambiguities and conflicts that made the dialogue embarrassing or threatening. Then consequences were bypassed by advice to be positive and patient. Indeed, such advice required that the bypass be covered up. The bypass strategy employed by the faculty at Crotonville was similar to the one used by the general managers in their back-home setting. This, we suggest, is why the Crotonville program and many other similar programs were dubbed "charm schools."

There was another problem that spurred the advice to "show your superiors that the new rules do give you the authority you want." There were general managers who took this advice seriously and pushed back against their superiors. Their superiors, in turn, tended to see such behavior as problematic. If the general managers met their goals for financial performance, their problematic behavior was discounted; they were seen as "diamonds in the rough." But those general managers who pushed back and did not meet their performance goals were eventually transferred, demoted, or fired.

From the experiences described earlier, we suggest the following conclusion: whenever double-loop changes in organizational theories-in-use are coupled with reasoning and interpersonal theories-in-use for organizational inquiry that remain consistent with Model I, self-fueling defenses that limit learning, along with such accompanying consequences as distancing and cynicism, will be activated whenever issues are embarrassing or threatening.

Under such conditions, it makes sense for managers to operationalize the social virtue of trust as "Just leave me alone, and monitor my results." Indeed, superiors were often proud to say to their subordinates that since they trusted them, they would leave them alone and simply monitor their results. Subordinates held a complementary concept of trust, that is, to distance themselves from their superiors. This eventually led to superiors being excluded from dialogue over business issues, even when they had the relevant knowledge and skills. Similarly, superiors excluded those subordinates with whose performance they were increasingly disappointed, thereby undermining their chances of promotion to the upper echelons of management.

Interestingly, there were some occasions when superiors and subordinates violated this concept of trust. These occasions were related to the possible activation of organizational defensive routines. Meetings about how to get around these routines would be held. Almost no meetings were held on how to eliminate or reduce them. One reason was that this possibility had never been legitimized during the early processes of implementation and education. Such a topic was treated as undiscussable, and its undiscussability was also undiscussable. A second reason was that no skills that might lead to double-loop changes in these phenomena were taught. A third was that organizational defensive routines were often seen as unfortunate but predictable features of all organizations.

Under these conditions, autonomy and authority could develop for issues that did not activate organizational defensive routines. When the defensive routines were activated, individuals often used their delegated authority to cover up their bypass activities. For example, in one large business unit, the senior executives approved spending two million dollars a month because they felt a certain project was close to the CEO's heart and he could not be influenced about the project. As several of these executives reported, it made little sense to waste precious "chits" on this issue.

As a consequence of bypassing situations of tension and ambiguity in the delegation of authority and of the hands-off attitude that resulted, it took a while before upper levels of management at GE discovered that profit center managers were attempting to meet bottom-line targets by dubious methods. Expected sales might be treated as actual, or reports of

losses might be delayed until the next quarter. Upper management responded to such discoveries by introducing tighter controls and information systems.

These consequences are more general. Whenever bypasses and cover-ups are discovered, the usual remedial action is to tighten controls, usually by including more detailed specifications. The difficulty with this strategy is that as the specifications mount, they:

- 1 violate the idea of profit center autonomy and responsibility; and
- 2 become so cumbersome that they tend to fall of their own weight.

We conclude, then, that as organizations move from centralized to decentralized policies and practices, a characteristic sequence develops over time:

- 1 Superiors and subordinates tend to report satisfaction and enhanced performance because of their newly delegated decision rights. A minority may resist taking charge because they are doubtful that top management "really means it" and/or because they are not yet confident of their own abilities.
- 2 As time goes by, issues surface over which strong differences in views come into play and cause conflicts. These consequences activate the participants' Model I theories-in-use and the organizational defensive routines that accompany them. They lead to superiors taking charge, thereby appearing, in the view of subordinates, to violate the concept of decentralization. The subordinates, in turn, may distance themselves and cover up. This, in the eyes of superiors, violates the responsibility they must assume if decentralization is to be effective. How often have we heard superiors warn subordinates that there must be "no surprises"? This warning applies not only to performance results but also to the use of organizational defensive routines.
- 3 Model I social virtues become strengthened. This, in turn, strengthens the defensive routines at all levels of the organization. For example, managers help and support by telling others what they believe will be positive and, therefore, reduce their hurt feelings. Managers show respect by not confronting the reasoning processes other managers use. Individuals who strive to be strong (advocating their positions in order to win), to be honest (not to lie), and to show integrity ("stick to one's principles"), will likely do so around issues that are relatively objective or where the control procedures make such action transparent. Whenever issues go underground and become undiscussable, Model I views of strength, honesty, and integrity lead to polarized positions that inhibit double-loop learning at the interpersonal level.
- 4 A perverse sense of trust – namely, trust coupled with mutual distancing – develops, making that distancing undiscussable.
- 5 This produces a space of free movement that allows behavior that violates organizational norms to emerge – specifically, ways of achieving results that meet the letter but violate the spirit of the policy.
- 6 When upper-level managers become aware of these consequences, they tighten controls, in turn violating the concept of decentralization.

Since Cordiner's time, advances in information technology have made it possible to develop tighter and more efficient controls and bring to the surface actions that would otherwise remain underground. As the rigidities of the self-fueling, antilearning processes become increasingly difficult to hide, upper-level managers are beginning to realize that managerial practices, if they are to succeed, require reeducation that includes a focus on individuals' Model I theories-in-use. It is at this point that organizational theories-in-use and interpersonal theories-in-use become interdependent.

It is ironic to note that when the connection between the organizational and interpersonal theories-in-use was not made, the reeducational programs became known as "charm schools." Now, when the connection is seen as more crucial than before, one often hears, "This is too difficult!" or "Is it really necessary?" or "Isn't this dangerous?" In the early days, the superficiality inherent in an incomplete approach to reeducation not only led to limited learning, it was also used by superiors and subordinates to cover themselves. Today superiors and subordinates may seek to protect themselves by raising the specter of "danger" and "Pandora's box." We believe that if concepts such as decentralization, empowerment, and personal responsibility are to preserve and encourage continuous improvement, the connection between the organizational and the interpersonal will have to be made.

Activity Costing

The following is an account of Robert Kaplan's attempts to introduce firms to activity-based costing (ABC), an information system aimed at improving organizational performance by enabling more precise attribution of costs and profitability to product lines (Argyris and Kaplan, 1994). Kaplan and his colleagues have developed an explicit structure for sponsoring and implementing an ABC project. Typically they begin with an activity-based cost model of organizational expenses and profitability which often reveals "many unexpectedly higher cost activities, processes, products, and customers" (p. 7) and proceeds to an action phase in which "management acts on the insights revealed by the ABC model to produce improved organizational performance" (p. 8). Kaplan and his colleagues distinguish key roles in this process: advocates for change, sponsors who have the ability to approve the required changes, change agents who take on the task of proving the concept and developing the more accurate information it requires, and targets, "the person or group whose behavior . . . is expected to change based on the newly revealed information" (p. 8).

Kaplan and his colleagues learned that they would regularly encounter resistance to action on improved information made available by ABC. Successful introduction of the system in one division of a firm was often followed by failure to disseminate it throughout the firm. One of Kaplan's manager-clients told him:

I'm not convinced yet that the organization is geared to making difficult decisions based on information. It took us five years to drop a product line that we knew was not making money for us. If we're not prepared to take tough decisions, we don't need a fancy new information system.

In response to such resistance, Kaplan worked with some of his clients to develop "organizational enablers," such as performance measures keyed to incentives, formation of task forces and introduction of group training exercises. But these enablers often proved insufficient to elicit managers' internal commitment to the changes necessary for successful implementation of ABC. Subtle forces still exist inhibiting the acceptance of new ideas whenever technical theories challenge existing organizational norms and long-established ways of doing business, thereby threatening or embarrassing "people who have much of their professional and self-feelings identified with decisions, actions, and organizational structures implemented in the past" (p. 13). For example, the analysis may reveal that "certain favored product lines or customers are highly unprofitable," and product managers responsible for introducing and maintaining these lines or the account managers responsible for the highly unprofitable customers "become threatened by the quantitative and defens-

ible evidence of their value-destroying activities” (p. 16). Or managers may be embarrassed “by the explicit recognition of excess capacity in areas under their authority,” or more generally, by the recognition that their traditional cost and performance measurement systems have “led to bad decisions . . . and a lack of focus on wasteful activities” (p. 17). The experience of threat or embarrassment typically gives rise to defensive reasoning, as illustrated by the following.

- It is better to fail with existing procedures than to fail trying something new, especially where business is okay and no one is really up against the wall. Believe me, I know.
- People are more used to managing expenses than the behavior that drives the expenses. That’s the reality of organizational life.

When champions of ABC seek to overcome such defenses, they often adopt strategies of change that exacerbate the defenses. In one instance, a top manager bypassed divisional managers, going directly to the factory to develop the ABC model, without seeking divisional managers’ approval or precommitment to the changes that would be required once the large number of unprofitable products were revealed. He avoided early confrontation with these managers by concentrating on building a defensible model of operations, inoculating himself against the resistance he anticipated by asking the controller and data processing manager to make an independent confirmation of the analysis. Ultimately, the division’s executives still rejected the analysis. And by framing the debate as “a conflict between an analytic, rational approach versus the executives’ ignorance and resistance to change, he provided himself with an excuse that absolved him from blame for the lack of effective action,” which he had anticipated (p. 17).

Argyris and Kaplan conclude their analysis by observing:

In every change that involves the introduction of a new technical theory of action whose correct implementation could be embarrassing or threatening, the advocates, sponsors, change agents, and targets who use the generic implementation approach will still face the challenge of dealing with the defensive reasoning of organizational participants, including themselves. (p. 18)

Studies of TQM

Total quality management (TQM), introduced in Japan by W. Edwards Deming just after World War II, has had a 50-year history. Its transfer to the United States began in earnest in the early 1980s, following American management’s awakening to the fact that Japan was outperforming US industry not only in costs but in quality and that Japan’s superior performance was due, in no small measure, to its practice of continuous quality improvement based on Deming’s theory and methods. Throughout the 1980s, TQM grew steadily in its scope of adoption and prestige. By the early 1990s, however, its popularity had peaked, and what Beer (1994) describes as a “literature of disillusionment” had begun to appear.⁴ We discuss two examples of recent essays in which the authors reflect on the mixed record of TQM interventions, the nature of the resistance TQM has provoked, and the limits of the TQM methodology.

Schneiderman (1992) writes as a professional TQM advocate and teacher, drawing on long experience in teaching organizations to implement TQM. His article is addressed to the naysayers, but it also reveals a thoughtful analysis of TQM’s limits. He begins with a now familiar description of TQM which includes three main elements:

- 1 identification and flowcharting of component organizational processes, such as manufacturing, product development, filing, and telephone answering;
- 2 framing each such process as a partnership of supplier and customer and accompanying it with explicit metrics and standards for its performance; and
- 3 involving all workers engaged in a process in continuous, unrelenting process improvement through the identification and elimination of defects.

Next, Schneiderman reports on his study of “nearly 100 cases of exemplary use of the PDCA (Plan-Do-Check-Act) cycle.” These are some of his main findings:

- The rate of improvement was constant, which Schneiderman interprets in terms of the concept of “half-life.” A half-life of six months for a process improvement activity means that there is a 50 percent reduction in defect level every six months (p. 6).
- As process complexity increased, the rate of continuous improvement declined. This was especially true where the “complexity” was “organizational.” “With multiple managers who are not part of the team and often have conflicting goals, the improvement process slows” (p. 11).
- The complexity of processes most in need of improvement in an organization increases with organizational level. Success of the quality movement in Japan led to “the natural desire to move improvement activity to the ranks of management,” but “QC Circles were trained to avoid, if possible, problems that involved human behavior, particularly when the behavior was that of others outside of the team, and to focus instead on machines, materials, and methods” (p. 21).
- TQM tools applied to complex management problems may produce a “false sense of rigor.” They seem to “cut through the ambiguity that surrounds complex problems. But the ambiguity is real; it is the forced solution that is unreal” (p. 27). Managers like other workers “repeatedly make good decisions in complex situations based on their experience or gut feel . . . [Their solutions] are often better than the descriptions that we force-fit into our analytical frameworks” (p. 27). But managers are often caught in dilemmas they find unacceptable when their intuitions put them at odds with the consensual conclusion that has been reached by a TQM team.
- A more fundamental limit of TQM, as practiced today, is that it fails to address the dynamic interactions that exist in complex systems, interactions that make it difficult to answer questions about the relative importance of multiple causes of defects in a process. Top management struggles with critical processes in which causal linkages are at best obscure. Consider the proposition that “increasing compensation will reduce labor turnover.” What if job satisfaction, not compensation, is the real issue, and employees, not wanting to burn their bridges, camouflage the real causes of unwanted attrition? Or consider the proposition that “centralization improves efficiency.” Why, then, do most reorganizations follow a pendulum that swings between centralization and decentralization?

Schneiderman concludes that TQM, as presently practiced, is limited. He even leans toward Ackoff's (1981) argument that “most current processes are so outdated and incapable of being improved in a timely way that they should be scrapped and replaced with idealized redesigns.” Nevertheless, he believes “the improvement process remains the same” (p. 35). He states that cross-functional problems of medium complexity can be remedied, in some cases, through the use of advanced TQM tools (the “7-management tools”), though “it is important to use them thoughtfully,” and “we must remember that consensus doesn't always lead to the best answer.” Finally, he suggests that highly complex,

interactive, dynamic problems which "seem to create the greatest amount of pain both inside and outside our organizations," might yield to the increased use of simulation modeling (p. 39).

Mallinger (1993) describes his experience as "change agent" in a company that proved relatively ineffective in its attempts to implement TQM. He spent a one-year sabbatical as participant/observer of a manufacturing division in a multibillion dollar Fortune 500 company. The corporate CEO, in response to financial and schedule pressures, announced a reorganization that drastically reduced management levels from 9 to 5 and cut management personnel by about 2000 people. The CEO then ordered the manufacturing division to implement TQM.

After five months of training and attempted implementation and an expenditure of \$11 million, the TQM program was halted, Mallinger reports, because of "a noticeable lack of transfer of TQM skills presented in training to the rework and delivery schedule problems faced by the company" (p. 11).

A year later, with a much-reduced training staff, the TQM intervention was begun again on the basis of an iterative approach to data gathering and intervention design. This time, a "lack of tribal knowledge" held by experienced managers, due in part to the lay-offs, led to a rise in skepticism about the process. Mallinger reports, based on numerous discussions with the division vice presidents, that TQM was not a priority and was seen, indeed, as an obstacle. "Although they each acknowledged the need for change, their actual behavior was quite different from their stated objectives" (p. 15). In one case a division vice president who indicated a desire to incorporate TQM as one of his five department goals for the year left a meeting after telling his human resource liaison staff to work out the plan to "fix my people." Apparently, the vice president saw TQM as a set of training techniques to improve the productivity of first-line managers rather than as a "commitment to mutual, proactive team planning." This individual seemed to think he would meet his obligation to be a TQM department, as ordered by the corporate CEO, by head count.

Middle managers also expressed skepticism and distrust, as indicated by comments such as, "Programs come and go around here, as do the people who support them." Shop floor people became increasingly discontented when the lay-offs began and advertised changes failed to occur. They felt their union leaders were joining their managers in taking them "down the road to hell." In one department, upper-level managers who had agreed to become active participants in the transformation process, dropped out of training meetings after the first hour of a twenty-four hour training program, claiming they had emergency meetings to attend. The first-level supervisors inferred that TQM had low priority for senior management and that they themselves had little value for the organization. They increased their skepticism of TQM and their distrust of their senior bosses. "No significant changes in performance within the business unit were reported two weeks, one month, or five months after training, and the intervention was curtailed" (p. 19).

As job insecurity increased with continuing lay-offs, the TQM team were increasingly seen as "hitmen on the prowl for fodder for the next round of reductions-in-force."

Mallinger reports further that he found no evidence that a "risk-taking, empowered environment" was in place. On the contrary, he found a pervasive climate of fear. This was illustrated by a meeting that the TQM team head, the human resources general manager, and the human resources vice president held with the president to brief him on the progress of the Phase II TQM model. The primary agenda was to sell the TQM design to the president rather than address issues of resistance such as those just described. Mallinger states:

The meeting took on a "we've got the answer now" message. I observed both the VP and GM agreeing with the president's philosophy and responding to each of his questions regarding

TQM in a way that suggested that we've got the "bases covered." It was evident to me that the project leader, who made the presentation, was uncomfortable raising substantive issues for fear of embarrassing his bosses. (p. 23)

Mallinger's report of his own behavior at the meeting is also of interest. He states that he "felt a responsibility not to collaborate in the ruse" but was also careful not to discomfort the HR managers. In hindsight, though, he felt he should have been more direct. What he did was to:

indicate to the president that the culture may be resistant to the changes that TQM advocates. His response was immediate: "First you change the system, then the culture will follow." ... The GM strongly agreed with the president, and the discussion of organizational culture ended." (p. 24)

The underlying meaning of what took place in the meeting, as Mallinger saw it, was to "not bring bad news to the President," a pattern he later observed at all levels of management.

The lessons Mallinger draws from this example are several. They include the need to form partnerships between the business unit under study and the elements outside it, to increase organization members' ownership of the TQM process by giving them greater input into its planning, and to encourage subordinates to engage in open discussion of troublesome issues in TQM implementation and in open expression of their fears and hopes for the process. However, Mallinger stops short of suggesting how, in corporate cultures like the one he observed, such lessons might actually be put into practice.

Roth's Study of Business Process Reengineering in Business Service Companies

George Roth's doctoral dissertation, written at MIT's Sloan School of Management (Roth, 1993), is a detailed analysis of attempts in two business service companies to improve productivity and quality of service through interventions that combined the introduction of new information technology with "business process reengineering". Roth defines this term, following CSC Index Inc., as "the concept of fundamentally changing the way work is performed in order to achieve radical performance improvements in speed, cost, and quality."

Roth undertook a year-long study of the introduction of imaging systems coupled with process reengineering in two insurance company subsidiaries. One of these, which he calls Dover Service Company, managed mutual fund investments for a financial services operation; the other, Harwick, provided service for life insurance policies. Both of these companies had similar work activities: processing paper-based information, answering customers' telephone inquiries, and performing business transactions on computer systems. Both were engaged in reengineering customer service and paper-processing operations through the development of imaging system applications which were intended to take the place of paper processing and were based on the same image system technology.

Roth organizes both of his studies around a stage model of the reengineering process that begins with strategy formulation and proceeds through phases of systems selection, preparation, implementation, and observation of use and outcomes. As he traces the development of the two projects through these stages, he describes the roles, activities, and views of four stakeholder groups:

- 1 *executives* (CEOs, presidents, COOs, and senior vice presidents), who had a role in formulating strategy and making decisions about imaging systems;

- 2 *managers*, who had line responsibility for daily business operations,
- 3 *technologists*, who were responsible for systems specification, development, installation, and operation, and
- 4 *workers*, the users of the imaging technology, who processed documents, responded to phone calls, and serviced customer requests.

Although the contexts of the two companies and their patterns of systems development and implementation are somewhat different, Roth's findings in the two cases are basically similar.

He finds in both studies that the anticipated benefits in service quality, speed, and productivity did not materialize. At Dover he writes:

A year later in May 1993, there were many changes but little effect on technological change outcomes. Use of the new account imaging application never resumed. Managers were still uncertain of where imaging system benefits would occur and how they would have a greater role in specifying changes. The Processing Department was allowed to hire only one person and assigned one manager part-time to work on imaging system developments. A new project to front-end scan all incoming documents in the mail room was undertaken and completed. This imaging application provided good statistics and management reports, but it did not produce productivity benefits, and documents continued to be processed in paper form. (p. 167)

At Harwick, Roth reports on outcomes across some eight sub-projects. Across the board, he writes, staffing levels in the two departments he focused on (Customer Service and Universal Life) were reduced by 23 percent, and "imaging system implementation effects on work, workers, and managers were not auspicious" (p. 244).

In Customer Service, Roth reports that one project, "UltraFind," designed to assist consultants in servicing customers over the phone, was found five months after implementation to provide capabilities required by "only one out of fifteen customer calls" (p. 249) and that "UltraFind was simply too slow to provide benefits for consultants' work" (p. 251). In the case of TIP, a paper processing system for company reps, Roth found that "reps required 3.3 times longer to process documents on TIP than to use paper documents and TAA (the preexisting system) (p. 256).

At Universal Life, an image processing application, ULIP (Universal Life Image Processing), replaced paper documents in workers' processing of cash items, franchise updates, loans, and the like. Roth reports that "workers' initial enthusiasm for ULIP waned when they were unable to process documents faster than their previous paper-based methods" (p. 261). Document processing time increased by a factor of 1.8, and questionnaire data showed "deterioration in workers' attitudes and work perceptions" (p. 265). Roth characterizes one application, CC (Correspondence Creation), designed to help word-processing workers by allowing them to select form letters from a menu, as successful. Correspondence consultants showed a 26 percent increase in productivity, and their responses to open-ended questions were more positive than those of any other workers (p. 275). Roth points out, however, that these responses were the exception rather than the norm for workers and managers.

What is especially interesting is not so much that Roth found an absence of anticipated increases in measurable productivity, quality of service, and worker attitude and performance, outcomes that may or may not be typical of the consequences of most business process reengineering, but that he described the interactions among stakeholder groups in the reengineering processes to which he partly attributes the disappointing project outcomes. Roth states that:

The two organizations' inability to create conditions in which stakeholders could communicate with one another led to fundamental difficulties in managing technological change and unexpected negative consequences. Stakeholders in technological change processes were unable to reconcile underlying assumptions, avoid miscommunications, or create shared meaning for imaging systems and their organizational implications. As a result, the more powerful stakeholder groups' ideas and agendas were promoted through imaging system activities. Technologists dominated the technological change process. (p. 299)

Roth gives many examples of what he means by "fundamental difficulties" of interaction. At Dover, for example, a line manager said that "he did not trust technologists because they withheld information and were unwilling to listen to his suggestions" (p. 157). He found that the technologists could not be influenced. He wanted his own specialists to help refine and adapt the application, but he was unable to secure resources for them to do so. Another manager complained that he relied on the technologists because they had chosen the imaging system and knew how to implement it, but "they applied [the system] broadly across the business process, without regard for whether it was appropriate, or for what business changes were required" (p. 158). Still another manager said that when change issues surfaced, "there was no process in place for deciding them," meetings were unsatisfactory (because they were dominated by technologists), and "there was no dialogue outside of the meetings."

What is striking about both of Roth's case examples is that, in spite of the disappointing data reported by managers and workers and independently ascertained by him, the top executives declared their projects to be resounding successes. Dwyer, the COO at Harwick, who had been hired to make organizational improvements, told Roth that "although she was aware of some difficulties," she believed that business objectives had been met, and she touted the project's success to the industry at large. At Dover, Bryant, the president, had a similarly positive evaluation of his reengineering project. Managers believed that Bryant held this view because he was fed only positive information; the problems were kept from him. They also believed that Bryant colluded with the masking of negative information because "he did not create conditions of intense scrutiny and involvement among his staff" (p. 138). And when lower-level managers tried to blow the whistle on this reporting process, they were punished. One manager said that she was floored when she saw a report indicating that the system had "98 percent up time," since she had her own experience of system failure. The technologist she first approached told her that, if she had problems with the reporting, she should call the project manager. When she did so, she was told that "it was reported that way because Craig (the CEO) wanted it that way." This manager concluded that Bryant had only received information that the project was going well and that the imaging system's problems and its real impact were never discussed at Bryant's staff meetings. The manager eventually "withdrew her attention from the imaging project because 'we aren't getting anywhere with [it] and I have other things to do'" (p. 138).

Roth concludes his analysis as follows:

Technologists had enough power in both organizations to coerce workers and managers to accept and use applications, but not enough power to assure applications' success. Technologists did, however, wield enough power to report mostly successful results despite contrary evidence in outcomes. (p. 302)

Beer's Study of His SHRM Intervention at Alpha Medical

Michael Beer's account of his intervention in the Alpha Medical Corporation (Beer, 1994) is an unusually full and frank exposition of an attempt to develop, not a one-time change, but

an organization capable of the kind of “continuous learning” suited to “the kind of strategic change required by competitive forces” (p. 6). Beer deliberately sought to bring about the internalization of a continuing organizational learning process that would link “hard” and “soft” elements of the organization. In his view, such a process must be systemic and iterative; it must establish a partnership between upper and lower levels of the organization; and it must “find some means for making normally hidden organizational and managerial behavior transparent” (pp. 7–8). In short, Beer’s intervention was directly addressed to the kinds of behavioral resistances the authors, cited earlier, said they encountered when they undertook or observed attempts to undertake interventions based on decentralization and management empowerment, TQM, or reengineering. It can also be seen as an attempt to create what Schein (1992) calls a “learning culture.”

In Beer’s Strategic Human Resource Management (SHRM) intervention, he hoped to link strategic change to a continuing process of organizational learning in the sense that “experience in implementing the new [strategic] design [would be] used to redesign the organization to align with limitations in human capabilities, or further changes in the environment discovered along the way” (p. 8). Beer began with “strategic profiling,” and went on to assess “organizational ability to reframe and implement new strategy matched to changing business realities” (p. 9). These and later steps were conducted by an employees task force made up of “the unit’s best employees one or two levels below the top team.” This team operated within structures that were designed to promote the discussion of findings and issues that might otherwise remain undiscussable. In the early profiling meeting:

While management listens, the [task force] sits in the middle of the room and discusses their findings [from investigations of the organization’s ability to implement its new strategy], organized into themes. This “fishbowl” discussion, interrupted at the end of each theme for questions of clarification from the top management team, typically goes on for several hours. (p. 10)

The organizational diagnosis, jointly conducted by the task force and top management, went on to identify “key design levers and human resource policies within the unit, including the top team itself, which are at the root of the problems identified by the task force.” Subsequently, the participants worked out a vision, philosophy of management, organizational design, and action plan for implementing and managing strategic change (p. 11). Among the “barriers to implementation” identified in this process were several that Beer believes to have been “undoubtedly responsible” for the failure of such enabling interventions as total quality management, reengineering, and employee involvement, including unclear or conflicting strategic priorities; difficulties in how the top team works together, top-down management style, poor interfunctional/divisional coordination, poor up/down communication, and deficiencies in career development and management competencies.

Participants in the SHRM intervention reported that it “allowed us to discuss the undiscussable . . . [putting] things on the table that would have taken [us] years,” giving rise along the way to a sense that “we didn’t know what we were getting into” and anxious jokes about the messenger being shot. Beer believes that because the project’s data had been generated by employees whom the managers had chosen as the best, the managers accepted the data “as valid and relevant, despite the fact that many issues point to their own effectiveness as managers” (p. 15).

Over the longer term, however, the outcomes of this intervention were viewed as mixed. Some intended objectives were achieved; others were not. Beer’s surveys and interviews

showed the following pattern. There was high commitment to the process of change at the top of the organization but less at the bottom (p. 16), and there was only modest change in overcoming barriers and accomplishing strategic tasks. The most change was in top-team effectiveness and cross-functional coordination; the least, in employee perception that top management rewards and promotes employees on the basis of their skill in managing and leading others and on their willingness to raise difficult issues with higher management.

Clearly, Beer concludes, top management "had not created a partnership with employees and a sense of group responsibility for producing strategic results" (p. 18). SHRM remained invisible at lower levels of the organization and remained limited by "hierarchical assumptions about how decisions should be made" (p. 17). In all but a few instances, higher management did not review SHRM results and follow them up, mainly, Beer believes, because of "their own discomfort and lack of skill in discussing difficult issues with subordinates" (p. 19). Beer observes that such discussion:

requires interpersonal competence and management values that the process itself is designed to develop . . . Though SHRM successfully opens a window for open dialogue, it apparently does not develop interpersonal and organizational capability to sustain that dialogue. (p. 20)

Beer's final word on his intervention is that a competitive world requires organizational capacity for self-design, which calls for management processes and skills to "engage in an open, fact-based conversation" about "difficult and often painful" issues. Such skills "have not yet been invented," he believes. They require "interpersonal behavior regarded by many managers as an 'unnatural act'" (p. 22). It remains to be seen, he concludes, "whether these skills can be developed and/or whether organizations must obtain them through selection" (p. 23). In either case, Beer asserts that these are "core skills for an effective organization," whether managers and their consultants recognize them to be so, and "they will be in greater demand as organizational adaptability becomes important in a more competitive world" (p. 23).

LESSONS FROM THE TWO SETS OF STUDIES

The Two Gaps

We have examined two sets of studies: the first deals with episodes in the lives of business firms in which organizational learning or adaptation occurred "naturally;" the second deals with interventions designed around various types of enablers that were intended to enhance organizational capability for productive learning.

The first set of studies (Van de Ven and Polley, 1992; Burgelman, 1994) reveal a *gap of explanation*. In both of the organizations studied, there were mixed results; there was some evidence of productive organizational learning, coupled with evidence that the organization had made only a partial or delayed response to data that could and should have been interpreted as calling for a significant change of course. Dynamically conservative, these organizations tended to persist in their previously established patterns of thought and action, and in so persisting, failed to act on what some of their members knew. During critical periods of decision, managers exhibited some patterns of inquiry that were distorted, overly constrained, and ineffective. The researchers take notice of these patterns, but they stop short of emphasizing or exploring them.

The two studies provide incomplete explanations of their mixed results. They do not fully explain first-order errors, such as the firm's failure to bring managerial attention to

bear on signs of trouble that were known to insiders in the early stages of the TAP case, or in Burgelman's study, Intel's lagging response to the shift in the DRAM market and its competitive position in that market. They do not at all explain the second-order errors that underlie first-order errors, namely, how key managers remained individually and collectively unaware of the firm's lagging response to important mismatches between the assumptions that guided organizational action and the changing patterns of internal or external reality, and how the organization kept undiscussable the issues raised by knowledge already privately held by some of its members.

The second set of studies (of decentralization during the Cordiner regime at GE, activity-based costing, TQM, reengineering, and SHRM) reveals a *gap of implementation*. The studies show that mixed results also followed from attempts to introduce enablers of productive organizational learning, that is, techniques and information systems (ABC), new formal structures (decentralization and flattening, creation of inquiry units, as in SHRM), principles and procedures for the development of a learning culture (TQM, reengineering, and SHRM). These interventions resulted in only partial or temporary improvement in organizational performance. For example, Cordiner's introduction of decentralization at GE was subverted, over time, by organizational games of authority, deception, and control. The diffusion of activity-based costing was stymied by the defenses it stimulated on the part of managers who were threatened or embarrassed by the information it would bring to light. The study of TQM interventions showed a declining rate of improvement in performance where the issues involved problems of behavior, such as conflicting goals and interests held by multiple managers, especially at higher levels of the hierarchy.

The studies of TQM, reengineering, and SHRM document the gap of implementation. Some of them (Beer, Schneiderman) point directly to the importance of constrained and/or distorted organizational inquiry but offer no general theory to account for it. In some instances, the researchers suggest hypotheses to account for the observed results of the interventions they studied, but their attempts to explain mixed outcomes by reference to cognitive limitations, interest theory, management values, and political games are only partly satisfactory. These hypotheses do not fully explain the first-order errors the researchers have identified in their cases. And they do not at all explain the second-order errors *we* find in their cases: the organization's being unaware of its inability to act on what it knew or to detect and correct error.

We have argued that the mixed outcomes and the limits of productive learning revealed by both groups of studies are attributable to the persistence of Model O-I learning systems, reinforced by the Model I theories-in-use of individuals whose thought and action are partly constrained by these systems. Under conditions of embarrassment or threat, these O-I systems distort processes of organizational inquiry. They give rise to defensive routines that block inquiry, they make critically important issues and dilemmas undiscussable, they prevent organizations from acting on what some of their members know, and they reinforce organizational awareness of the ability to detect and correct error.

We turn now to a discussion of the implications of our analysis for the principal challenges to a theory of productive organizational learning. We explore how our theory-of-action account of the gaps of explanation and implementation bears on the several doubts raised by learning skeptics about organizational learning:

- its meaningfulness;
- its desirability;
- the capability of real-world organizations to engage in productive versions of it; and
- the efficacy of interventions aimed at enhancing it.

The Paradox of Organizational Learning

The meaning of “organizational learning” hinges, as we have seen, on the crucial issue of the levels of aggregation at which organizational phenomena are described and explained and at which prescriptions for organizational action are directed. Our analyses of the two groups of studies show that one cannot account for the observed higher-level phenomena of organizational learning, that is, those that seem important to researchers concerned with strategy making or technological innovation, without referring to individual and interpersonal processes of inquiry. The feedback loops contained in our cause maps and models of O-I learning systems show crucially important causal linkages among three levels of aggregation: interpersonal inquiry, interactions among organizational subunits, and the patterns of action and learning characteristic of whole organizations. We have argued that we should also treat as organizational the kinds of processes that shape the phenomena we take to be organizational in the first instance.

Unless we refer to the level of interpersonal inquiry, we cannot:

- explain the empirically discovered phenomena of limited organizational learning;
- resolve plausible but conflicting explanatory hypotheses about limited organizational learning;
- produce an explanation useful for redesigning practices so as to overcome barriers to productive organizational learning.

The Meaning of “Productive” Organizational Learning

What do these examples of mixed results tell us about the nature of productive, especially double-loop, organizational learning?

They indicate the occurrence of some double-loop changes in organizational outcomes at the level of theory of action. But these changes are only temporary (one-shot) or affect only a part of the organizational theory of action and not others. In both groups of examples, the organization’s theory-in-use for organizational inquiry did not undergo double-loop change. Although in Beer’s case the intervenor did seek to inject Model II values into organizational inquiry, the injection did not last beyond the period of intervention.

It is clear, then, that it is possible for limited double-loop learning to occur at the level of outcomes in organizational theory of action without double-loop learning in processes of organizational inquiry. We have argued that whatever values may be espoused for organizational learning – single-loop, instrumental improvement in organizational performance, or double-loop change in the criteria by which organizational improvement is measured – their sustained achievement depends on the organization’s continuing ability to engage in double-loop organizational inquiry. The Model II values that govern double-loop organizational inquiry are essential to the sustained achievement of other criteria for productive learning, for example, sustained and timely adaptive learning in technological innovation (Van de Ven and Polley), the timely reconciliation of formal strategy with a changing strategic situation (Burgelman), or the maintenance of the productive effects of enablers, such as decentralization and management empowerment (as in Cordiner at GE), or of structures for enhanced organizational inquiry (as in Beer’s SHRM intervention).

The Model II values that govern double-loop organizational inquiry are foundational to sustained productive organizational learning. By the same token, as we have shown in each of the studies we have discussed, double-loop organizational inquiry with its

Model II values is required in order to disrupt the defensive routines that preserve the organizational status quo in situations of ambiguity or uncertainty by keeping threatening issues undiscussable and by keeping participants unaware of how they help to perpetuate those routines.

Our emphasis on Model II values for double-loop inquiry does not obviate the issue of learning for evil ends. In any given instance, the objectives of organizational learning and the values that inform those objectives must be examined. We would argue, however, that double-loop inquiry, which seeks to illuminate issues and dilemmas we might otherwise suppress or of which we might otherwise remain unaware, offers the most likely route to such examination.

Threats to Productive Organizational Learning

In chapter 10 we will discuss three kinds of threats: to coherent organizational action, to cognitive capability for learning, and to effective implementation of lessons drawn from experience.

The issue of coherence arises in several of the cases we analysed in this chapter, most obviously in Van de Ven and Polley's and Burgelman's cases. In all of these cases, there was evidence of incoherence in relationships between middle and top management, as well as among divisions at the level of middle management. Top and middle managers or divisions such as R&D and Fabrication, diverged from one another, enacted contrary policies, or engaged in a kind of parallel play, pursuing independent policies without coordination. Burgelman now argues that in the early stages of his example this lack of coherence proved actually to be useful since it allowed DRAM development to go on beyond the point when at least some people in the company clearly saw that Intel no longer had a competitive edge in the DRAM business. The result was that the company could eventually take advantage of the unanticipated benefits of that development for the microprocessor business that it was (by then) coming to see as the wave of the future. In this case, if we accept Burgelman's analysis, a vein of technological development undertaken for the sake of one business proved unexpectedly to be useful for another business. Such things happen. But the outcome might have been less rosy. At crucial moments in Van de Ven and Polley's story of the TAP development and in Roth's story of business processes reengineering at Dover and Harwick, analogous situations arose when TAP's development and the reengineering processes continued beyond the points at which they should have been subjected to critical scrutiny. In these instances, according to the authors, the results were negative.

What lessons should we draw from such occurrences? We certainly should not conclude that all programs of technological development undertaken for the sake of a business now seen as obsolescent should be continued because they might turn out to be useful for another business. On what basis, then, should one distinguish projects worth supporting from those that ought not to be supported? Even to raise this question is to pose the issue of awareness. If organizational inquirers are not aware of an existing incoherence in policies enacted by different levels or groups of managers and if they are not capable of subjecting that incoherence to public scrutiny, then how could such discrimination be made? In both the Van de Ven and Polley and Burgelman cases, managerial awareness and public discussion of disjointed policies would have required double-loop inquiry into the issues that provoked threat or embarrassment and triggered organizational defensive routines – inquiry comparable to the inquiry undertaken by the CEO and directors of the management consulting firm described in chapter 9. The incoherence that might be taken as evidence of lack of capability for productive learning in real-world organizations, giving rise to generalizations of the kind

that March and like-minded scholars have proposed, could also be seen – and is in fact seen, from the theory-of-action perspective – as an occasion for double-loop inquiry.

All of the cases we have considered in this chapter provide evidence that could be attributed to *limited cognitive capability for productive organizational learning*. The question is how we ought to interpret this evidence.

In several of these cases, threats to valid inference arose. In most of them, uncertainty and ambiguity surrounded the causes of some critically important effects of organizational action. In some instances (Intel), managers faced with uncertainty fell into competence traps or engaged in superstitious learning. In other instances (TAP), managers failed to focus concerted attention on signs of trouble that ought to have cast doubt on the assumptions that underlay established organizational policies. In most of the cases of attempts to introduce enablers to improve organizational capability for productive learning, managers persisted in a course of action about which some members of the organization had significant grounds for doubt, yet the information and interpretations held by these managers did not give rise to a concerted, public critique of the assumptions that underlay organizational action.

Should we argue, as March and others have done, that such instances of first-order error provide further evidence for the general proposition that organizations lack cognitive capability for productive learning from experience? Our position is that such instances do not, of themselves, reveal limited cognitive capability. Rather, a close inspection of the evidence compiled by the authors of the case studies reveals a pattern of organizational defenses that reinforce a prevailing blindness to first-order error, a blindness that makes such errors uncorrectable and their underlying, mistaken assumptions self-sealing. The question of cognitive capability must remain moot since during the time periods in question, the fact that organization members were unaware of or unwilling to discuss the relevant issues, prevented the organization from applying to them whatever cognitive capability it had. There is, indeed, some evidence, from the authors' account of patterns of inquiry carried out in the later stages of the Intel and the TAP cases, that both organizations possessed the ability to engage in complex and refined exercises of error detection, causal analysis, and error correction.

There is, of course, no guarantee that if the managers in these cases had been aware of their first-order errors and had subjected them to public discussion and critical analysis, they would have succeeded in correcting them. In the world of management, as of science, there is no guarantee against persistent error. But there is also no a priori reason to attribute persistent first-order error to the limits of cognitive capability *when there is manifest evidence that a prevailing pattern of defensive behavior blocks the exercise of that capability, whatever its strengths or limitations may be*. The much more important issue raised by our analysis of the case material is the need and potential for double-loop inquiry into the patterns of second-order error, the limited learning systems whose persistence makes first-order error unassailable. This issue is central to the gap of implementation and critical to the design of interventions aimed at enabling productive organizational learning.

Moreover, from the perspective of scholarly research on organizational learning, a robust test of propositions about the cognitive capability of organizations depends on making organizational actors aware of, and helping them to overcome, the defenses that blind them to first-order errors. For in the absence of such an intervention, how could one tell whether an organization's vulnerability to such threats to valid inference as competence traps and superstitious learning is due to the organization's limited cognitive capability or to the effects of its organizational defensive routines?

Conceivably, researchers might argue along the lines of the behavioral theory of the firm that defensive routines are unalterable facts of organizational life that should be "folded

into" the attribution of limited cognitive capability. We would agree that patterns of organizational defensiveness are widespread and deeply embedded in the behavioral worlds of organizations. But an intervention of the kind reported in chapter 9 and others like it (see Argyris, 1990, 1993), stand as existence proofs of the possibility of double-loop inquiry that breaks through the constraints of O-I behavioral worlds. In the wide-ranging universe of organizational life, such episodes must be counted as rare events. Nevertheless, they cast doubt on the validity of any general attribution of limited cognitive capability and point the way to future programs of intervention and research.

Overcoming Impediments to Productive Organizational Learning

Our analysis of case material related to decentralization, TQM, activity-based costing, reengineering, and SHRM, indicates that organizational enablers can induce productive organizational learning, leading even to temporary double-loop outcomes at the level of organizational theory-in-use, when the issues are neither embarrassing nor threatening. But under conditions of embarrassment or threat, we find that the O-I learning systems of the organizations tend over time to subvert the enabling interventions.

We believe that our analysis helps to explain the characteristic life cycles of such organizational fixes as TQM, flat organization, reengineering, and management empowerment. In each such instance, a prescription for organizational reform appears on the horizon, supported by plausible-sounding theory and stories of successful implementation by early adopters. Often at the core of the reform there lies a significant insight, for example, managers *should* be freed up to take on greater responsibility and make greater use of local knowledge, or organizational processes *should* be rethought in the light of the possibilities opened up by advanced information technology. Usually, however, the prescription is converted by its advocates and the consultants who undertake its dissemination into a readily understandable package of procedures. Not infrequently, the package is accompanied by an ideology that takes on quasi-religious overtones. Organization managers, thirsty for solutions to the persistent predicaments in which they find themselves and impatient with calls to wrestle with the complexity of the predicaments or with their own possible collusion in reinforcing them, latch onto the package. A bandwagon effect ensues as managers adopt the package because managers around them are adopting it. Then over time, as experience with the reform builds up and as good intentions are subverted by organizational defensive routines and their associated Model I theories-in-use, a literature of disillusionment begins to appear. Then normal cynicism begins to reassert itself. Lower-level managers begin to mutter that those at the top never really meant it, and top-level managers express frustration at the intractable resistance to change exhibited by those below them. People begin to say, "We tried that!" and a readiness for the next reform package begins to take shape.

Nevertheless, as we have argued, there are examples of rare events in which individual members of organizations show themselves willing and able to engage the complexity of the limited learning systems in which they operate, to examine their own skilled contributions to the maintenance of those systems, and to learn to move toward the Model II values, strategies, and assumptions that enable them to disrupt prevailing organizational defensive routines. Perhaps as Beer and others have suggested, some managers who have not experienced such rare events but have experienced the reforms described in literatures of disillusionment are growing in readiness to wrestle with complexity and to explore the patterns of thought and action through which they contribute to their own disillusionment.

These conclusions hold important implications for practice and research.

Practitioners who want to increase their organization's capability for continuing productive learning, especially of the double-loop variety, should learn to improve the performance of organizational inquiry, which requires double-loop learning in their own theories-in-use.

Researchers who want only to explain the phenomena of limited organizational learning should attend to the directly observable data of organizational inquiry which must, in turn, be linked to higher levels of aggregation. This strategy of attention is all the more appropriate for researchers who want to produce usable knowledge that can help practitioners build organizational capability for productive learning. In both cases, researchers would do well to reframe their research as action research undertaken in collaboration with practitioners who seek to build capability for enhancing and expanding the rare events of productive organizational learning.

Notes

- 1 Draws heavily from Chris Argyris and Donald Schön (1996).
- 2 In order to test their hypotheses, the authors decompose their case into "events," which they define as "critical incidents when changes were observed to occur in the innovation idea and activities, personnel appointments and roles, innovation unit relationships with others, environmental and organizational context, and outcomes" (p. 99). They adopt a monthly interval for the aggregation of events, collecting some 325 of them over a five-year period. They code outcomes of events as positive or negative or mixed, depending on their embodying good news or successful accomplishments, bad news or instances of failures or mistakes, or combinations of these. They seek to test their hypotheses by ascertaining correlations or lack of them, between positive, negative, and mixed events, on the one hand, and, on the other hand, continuation or change in a prior action course, resource controllers' interventions, shifts in outcome criteria, or events (judged significant) in the external context.
- 3 Burgelman's argument here is reminiscent of Hirschman's famous "principle of the Hiding Hand." Hirschman (1958) proposed that ignorance of dangers to come can actually foster the success of a development project because those who underestimate the dangers also tend to underestimate the creativity that may be brought to their resolution. However, in his later writings, Hirschman also recognized the difficulty of distinguishing instances in which the Hiding Hand proves to be beneficent from those instances in which it does not prove to be beneficent!
- 4 See Michael Beer, 1994. On page 2, Beer writes that "70% of all corporations report that TQM has not lived up to their expectations."

References

- Ackoff, R., 1981, *Creating the Corporate Future*, New York, John Wiley and Sons.
- Argyris, C., 1990 *Overcoming Organizational Defenses*, Needham, MA, Allyn-Bacon.
- Argyris, C., 1993, *Knowledge for Action: A Guide to Overcoming Barriers to Organizational Change*, San Francisco, Jossey-Bass.
- Argyris, C., and Kaplan, R., 1994, "Implementing New Knowledge: The Case of Activity-Based Costing," *Accounting Horizons*, 8, 3, pp. 83-105.
- Argyris, C., and Schön, D., 1996, *Organizational Learning II*, Reading, MA., Addison-Wesley pp. 200-42.
- Beer, M., 1994, "Developing an Organization Capable of Implementing Strategy and Learning," Harvard Business School Working Paper, September.
- Burgelman, R., 1994, "Fading Memories: A Process Theory of Strategic Business Exit in Dynamic Environments," in *Administrative Sciences Quarterly*, March, 39; pp. 24-56.
- Hirschman, A. O., 1958, *The Strategy of Economic Development*, New Haven; Yale University Press.
- Mallinger, M., 1993, *Ambush Along the TQM Trail*, Mullica, Malibu, Calif., Pepperdine University.

- Ross, J., and Staw, B. M., 1986, "Expo 86: An Escalation Prototype", in *Administrative Science Quarterly*, 31, pp. 274-97.
- Roth, G., 1993, *Business Process Re-Engineering in Business Service Companies*, doctoral dissertation, Sloan School of Management, Massachusetts Institute of Technology, Cambridge, Mass.
- Schein, F., 1985, 1992, *Organizational Culture and Leadership*, San Francisco, Jossey-Bass.
- Schniderman, A. M., 1992, "Are There Limits to Total Quality Management?", mimeo, Analog Devices.
- Schön, D. A., 1967, *Technology and Change: The New Heraclitus*, New York, Delacorte Press.
- Van de Ven, A., and Polley, D., 1992, "Learning While Innovating", in *Organization Science*, 3, 1, February, pp. 92-115.