

organizational and its learning product an organizational one, if (eventually as happened with McLain) the covert project is discovered, publicly legitimized, and formally adopted by the organization.

There are instances in which organizational inquiry produces a temporary change in organizational theory-in-use, but the new understandings associated with that change, held only in the minds of certain individual "carriers," are lost to the organization when they leave. This often happens in small professional organizations, such as research and development firms, design offices, or software companies, where staff members habitually move in and out of organizational homes, taking their ideas and capabilities with them. Sometimes members of the host organization recognize the risk of loss and seek deliberately to uncover and document the special insights and skills of these organizational birds of passage.

Finally, it is worth noting that not all changes in organizational theory-in-use qualify as learning. For example, changes in an organization's environment (such as a slackening of product demand) may trigger patterns of response that undermine organizational norms. Members may lose enthusiasm, become sloppy in task performance, or lose touch with one another. These kinds of changes are forms of deterioration. On the other hand, as in an example we will describe in Chapter 9, an episode of organizational inquiry may provoke deterioration in some regions of organizational theory-in-use even as it generates learning in another.

All such near misses suggest that our category of organizational learning has fuzzy edges. There are boundary instances in which it is difficult to determine with precision whether interactive inquiry is truly organizational or has truly changed organizational theory-in-use or whether its results have been truly embedded in the organization's memories, maps, and programs. Such vagueness may be inherent in organizational phenomena that are ill-formed or emergent or may reflect, on the other hand, a lack of information sufficient to permit a clear determination. They do not invalidate our definition as long as many examples do clearly fall inside it and as long as what information we would need in order to make a definite attribution of organizational learning is clear.

Productive Organizational Learning

There are several ways in which instrumental learning may be for ill rather than for good. Some of these are particular to organizational learning; others, applicable to learning by agents of any kind.

First, the *ends* of action may be reprehensible. The value we attribute to an increase in effectiveness or efficiency depends on how we answer the question, Effectiveness or efficiency for what? and how we evaluate the "what." This issue is critically important when the action in question emanates from an organization whose members are eager or unthinkingly compliant participants. During World War II, Eichman's bureaucracy learned over time to become more efficient at sending its victims to the gas chambers.

The value attributed to a particular instance of learning also depends on how we judge its validity. Learning seems to suggest the acquisition of valid, workable knowledge or know-how. But when we treat organizational learning as inquiry that leads to a change in theory-in-use, we open up the possibility that any given change may be based on a lesson that turns out to be false or unworkable. James March (1988) uses the term "superstitious learning" to refer to one such class of lessons: those based on the belief that because events have followed one another in time they are also related to one another as cause to effect. For example, corporate managers' may believe that a rise in profits following the institution of a new policy must have been caused by that policy, though it may have been due to nothing more than an improvement in market conditions. March suggests that managers are drawn to superstitious learning because it reinforces the myth of managerial control—a belief congenial to the norms of managerial stewardship but often contrary to fact.

Organizational learning that is valid or workable at the time of its first occurrence may lead to effects that are negative overall. To take a notable example, "competence traps" (also March's term) are situations in which an experience of perceived success leads an organization to persist in a familiar pattern of thought and action beyond the time and conditions within which it yields successful outcomes. The behavior that yields success at time, *t*, may not yield it at *t*+1. Yet an organization lulled by its success and misguided by the lessons drawn from it, may persist in a familiar pattern of behavior long after it has ceased to work. In business strategy, General Motors, IBM, and Digital Equipment Corporation come to mind as recent examples of firms that persevered in following a once-winning strategy that had become a losing strategy, apparently blind to the fact that the competitive environment had shifted out from under it. Such examples should be understood in terms of the webs of interest organizations build up around familiar strategies, technologies, or structures, and the "dynamically conservative" processes (Schön, 1967) that

reinforce an organization's adherence to the lessons it has drawn from past experience.

Later we will have opportunity to see how people can learn collectively to maintain patterns of thought and action that *inhibit* productive organizational learning. For example, they may learn to respond to error by the use of scapegoating, games of unilateral control and avoidance of control, systematic patterns of deception, camouflage of intentions, and maintenance of taboos that keep critical issues undiscussable. Such patterns of thought and action, learned from experience, often have the effect of inhibiting the kinds of productive learning that yield improved performance or restructured values for performance. Yet members of the organization may develop an attachment to these patterns, even to the point of exclaiming, "It has taken us years to learn to live in this screwed-up world; don't make waves!"

If we were to use learning only in a positive sense, then we would have to qualify the learning involved in all such negative examples with adjectives like dysfunctional, pseudo, or limited. These semantic devices are misleading, since they tend to be applied to learning products after the fact; whereas we are often uncertain in any given situation of action, whether an alleged instance of productive organizational learning is valid and workable. The crucial point is that, as we try to understand or enhance organizational learning, we should keep in mind the variety of ways in which any particular example of it may prove to be invalid, unproductive, or even downright evil.

For these reasons, it is useful to distinguish three types of productive organizational learning:

1. *organizational inquiry*, instrumental learning that leads to improvement in the performance of organizational tasks;
2. inquiry through which an organization explores and restructures the values and criteria through which it defines what it means by improved performance; and
3. inquiry through which an organization enhances its capability for learning of types (1) or (2).

Single- and Double-Loop Learning

By *single-loop learning* we mean instrumental learning that changes strategies of action or assumptions underlying strategies in ways that leave the values of a theory of action unchanged. For example, qual-

ity control inspectors who identify a defective product may convey that information to production engineers, who, in turn, may change product specifications and production methods to correct the defect. Marketing managers, who observe that monthly sales have fallen below expectations, may inquire into the shortfall, seeking an interpretation they can use to devise new marketing strategies to bring the sales curve back on target. Line managers may respond to an increase in turnover of personnel by investigating sources of worker dissatisfaction, looking for factors they can influence, such as salary levels, fringe benefits, or job design, to improve the stability of their work force.

In such learning episodes, a single feed-back loop, mediated by organizational inquiry, connects detected error—that is, an outcome of action mismatched to expectations and, therefore, surprising—to organizational strategies of action and their underlying assumptions. These strategies or assumptions are modified, in turn, to keep organizational performance within the range set by existing organizational values and norms. The values and norms themselves (related in the previous examples to product quality, sales level, or work force stability) remain unchanged.

By *double-loop learning*, we mean learning that results in a change in the values of theory-in-use, as well as in its strategies and assumptions. The double loop refers to the two feedback loops that connect the observed effects of action with strategies and values served by strategies. Strategies and assumptions may change concurrently with, or as a consequence of, change in values.¹ Double-loop learning may be carried out by individuals, when their inquiry leads to change in the values of their theories-in-use or by organizations, when individuals inquire on behalf of an organization in such a way as to lead to change in the values of organizational theory-in-use.

¹ We borrow the distinction between single- and double-loop learning from W. Ross Ashby's *Design for a Brain* (New York: John Wiley and Sons, Inc., 1960). Ashby formulates his distinction in terms of (a) the adaptive behavior of a stable system, "the region of stability being the region of the phase space in which all the essential variables lie within their normal limits," and (b) a change in the value of an effective parameter, which changes the field within which the system seeks to maintain its stability. One of Ashby's examples is the behavior of a heating or cooling system governed by a thermostat. In an analogy to single-loop learning, the system changes the values of certain variables (for example, the opening or closing of an air valve) in order to keep temperature within the limits of a setting. Double-loop learning is analogous to the process by which a change in the setting induces the system to maintain temperature within the range specified by a new setting. See especially pp. 71-75.

Organizations continually engaged in transactions with their environments regularly carry out inquiry that takes the form of detection and correction of error. Single-loop learning is sufficient where error correction can proceed by changing organizational strategies and assumptions within a constant framework of values and norms for performance. It is instrumental and, therefore, concerned primarily with effectiveness: how best to achieve existing goals and objectives, keeping organizational performance within the range specified by existing values and norms. In some cases, however, the correction of error requires inquiry through which organizational values and norms themselves are modified, which is what we mean by organizational double-loop learning.

In any particular instance of double-loop learning, the resulting changes in values and norms may not be judged to be desirable: their desirability can be determined only through a situation-specific critique of the changes themselves and of the inquiry through which they are achieved. Nevertheless, it is through double-loop learning alone that individuals or organizations can address the desirability of the values and norms that govern their theories-in-use.

Consider a chemical firm which has set up a research and development division charged with the discovery and development of new technologies (an example we consider at greater length in Chapter 3). The firm has created its new R&D division in response to the perceived imperative for growth in sales and earnings and the belief that these are to be generated through internally managed technological innovation. However, the new division generates technologies that do not fit the corporation's familiar pattern of operations. In order to exploit some of these technologies, the corporation may have to turn from the production of intermediate materials, which it is familiar, to the manufacture and distribution of consumer products with which it is unfamiliar. This, in turn, requires that members of the corporation adopt new approaches to marketing, managing, and advertising; that they become accustomed to a much shorter product life cycle and to a more rapid cycle of changes in their pattern of activities; that they, in fact, change the very image of their business. And these requirements for change come into conflict with another sort of corporate norm, one that requires predictability in the management of corporate affairs.

Hence, the corporate managers find themselves confronted with conflicting requirements. If they conform to the imperative for growth, they must give up on the imperative for predictability. If they

decide to keep their patterns of operation constant, they must give up on the imperative for growth, insofar as that imperative is to be realized through internally generated technology. A process of change initiated with an eye to effectiveness under existing norms turns out to yield a conflict in the norms themselves.

If corporate managers are to engage this conflict, they must undertake a process of inquiry which is significantly different from the inquiry characteristic of single-loop learning. To begin, they must become aware of the conflict. They have set up a new division that has yielded unexpected outcomes; this is an error, a surprise. They must reflect upon this surprise to the point where they become aware that they cannot deal with it adequately by doing better what they already know how to do. They must become aware that they cannot correct the error by getting the new division to perform more efficiently under existing norms; the more efficient the new division is, the more its results will plunge the managers into uncertainty and conflict. The managers must discover that it is the norm for predictable management which they hold, perhaps tacitly, that conflicts with their wish to achieve corporate growth through technological innovation.

Then the managers must undertake an inquiry that resolves the conflicting requirements. The results of their inquiry will take the form of a restructuring of organizational norms and very likely a restructuring of strategies and assumptions associated with those norms; these must then be embedded in the images and maps that encode organizational theory-in-use. There is in this sort of episode a double feedback loop which connects the detection of error not only to strategies and assumptions of effective performance but to the values and norms that define effective performance.

In such an example of organizational double-loop learning, incompatible requirements in organizational theory-in-use are characteristically expressed through a conflict among members and groups of members. One might say that the organization becomes a medium for translating incompatible requirements into interpersonal and intergroup conflict.

For example, some managers of the chemical firm may become partisans of growth through research; while others, committed to familiar and predictable patterns of corporate operation, become opponents of the new, research-based conception of the business. Double-loop learning, if it occurs, will follow from the process of inquiry by which these groups of managers confront and resolve their

dispute. They may respond in several ways, not all of which meet the criteria for organizational double-loop learning.

First, the members may treat the conflict as a fight in which choices among competing requirements are to be made, and weightings and priorities are to be set on the basis of dominance. The R&D faction, for example, may include the chief executive who is able to win out over the old guard because of his greater power, or the two factions may fight it out to a draw, settling their differences in the end by a compromise that reflects nothing more than the inability of either faction to prevail over the other. In both of these cases, the conflict is settled for the time being but not by a process that could be appropriately described as learning. If the conflict ends with a power play or a stalemate, neither side is likely to emerge with a new sense of the nature of the conflict, its causes and consequences, or its meaning for organizational theory-in-use.

On the other hand, the adversaries may engage their conflict through inquiry in any of the following ways:

- a. They may invent new strategies of performance that circumvent the perceived incompatibility of requirements; they may succeed in defining a kind of research and development addressed solely to the existing patterns of business that offer the likelihood of achieving existing norms for growth. They will then have succeeded in finding a single-loop solution to what at first appeared a double-loop problem.
- b. They may carry out a trade-off analysis that enables them to conclude jointly that so many units of achievement of one norm are balanced by so many units of achievement of another norm. On this basis, they may decide that the prospects for R&D payoff are so slim that the R&D option should be abandoned, and with that abandonment there should be a lowering of corporate expectations for growth. Or they may decide to limit R&D targets so that the disruptions of patterns of business operation generated by R&D are confined to particular segments of the corporation. Here there is a compromise among competing requirements, but it is achieved through inquiry into the probabilities and values associated with options for action.
- c. The incompatible requirements may be perceived as incommensurable. In such a case, the conflict may still be resolved through inquiry that gets underneath the members' initial commitments. Participants must then ask why they

hold the positions they do and what the positions mean. They may ask what factors have led them to adopt particular standards for growth in sales and earnings, with what rationales, and what are likely to be the consequences of attempting to meet the standards by any means whatever. Similarly they may ask what kinds of predictability in operations are of greatest importance, to whom they are important, and what conditions make them important.

Inquiry of type B or C may lead to a restructuring of corporate values and norms. Or it may lead to the invention of new patterns of incentives, budgeting, and control that take greater account of requirements for both growth and predictability.

In this type of organizational double-loop learning, individual members resolve interpersonal and intergroup conflicts that express incompatible requirements for organizational performance. They do so through organizational inquiry that creates new understandings of the conflicting requirements—their sources, conditions, and consequences—and sets new priorities and weightings of norms, or re-frames the norms themselves, together with their associated strategies and assumptions. In such a process the restructured requirements for organizational performance become more nearly compatible and more susceptible to effective realization. And the resulting understandings, priorities, and reframed norms become inscribed in the images, maps, and programs of the organization and are thereby embedded in organizational memory.

Additional Considerations

The distinction between single- and double-loop learning is complicated by several factors. As we consider these, we identify gradients of significance in organizational learning, become aware of zones of ambiguity at the boundaries of these two types of learning, and identify a variety of patterns of inquiry through which organizations may engage in double-loop learning.

First, the distinction between single- and double-loop learning is complicated by organizational size and complexity.

Organizational theories-in-use are structures composed of many interconnected parts. Some of these are local and peripheral, while others are core elements fundamental to the structure as a whole. In a chemical firm, for example, norms governing requirements for growth and predictability are fundamental to the theory-in-use of the whole organization. If these norms were to change, a great

assumptions of the sort that practitioners ordinarily detect and try to correct. They tend to be selectively inattentive to second-order errors, which are due to the organizational designs that make people systematically unaware of the behavioral phenomena that underlie the production and reproduction of first-order errors. We refer here, for example, to defensive routines, mixed messages, taboos on the discussability of key issues, games of control and deception, and organizational camouflage. As we have argued throughout this book, reflection on such phenomena and the theories-in-use that underlie them, is essential both to the task of explaining the limitations of organizational learning and to the design of interventions that can overcome those limitations.

The Scholarly Literature of Organizational Learning

This literature—intentionally distant from practice, nonprescriptive, and value-neutral—focuses on just those questions the first branch ignores: What does “organizational learning” mean? How is organizational learning at all feasible? What kinds of organizational learning are desirable, and for whom and with what chance of actual occurrence? The scholars of organizational learning generally adopt a skeptical stance toward these questions. Their skepticism tends to revolve around three main challenges.

1. There are those who argue that the very idea of organizational learning is contradictory, paradoxical, or quite simply devoid of meaning.
2. A second challenge to the idea of organizational learning accepts it as a meaningful notion. What it denies is that organizational learning is always or ever beneficent.
3. A third kind of skepticism about organizational learning questions whether real-world organizations do learn proactively, and whether, in principle and in actuality, they are capable of coming to do so.

Organizational Learning Is Contradictory

As we stated in Chapter 1, when we begin by assuming that individuals are the only proper subjects of learning and that we know what we mean when we say that individuals learn, then we are likely to be puzzled and disturbed by the notion that learning may also be attributed to organizations. Indeed, some researchers have argued, as

Geoffrey Vickers did, that if the term, “organizational learning,” means anything, it means learning on the part of individuals who happen to function in an organizational setting. From this perspective, to say that an organization learns is to commit what the philosopher, Gilbert Ryle, called a “category mistake.”

Yet even a cursory reading of the recent literature suggests that the disposition to regard organizational learning as a paradoxical idea was far more vigorous twenty years ago than it is now. Economists such as Marin (1993), Herrstein (1991), and Holland and Miller (1991) have begun to introduce learning into economic discourse, making explicit the references to learning that have long been implicit in such branches of economics as the theory of the firm, which treats the firm as a decision maker optimizing to a utility function, and theories of free-market competition, which give an essential place to gains in efficiency and productivity stimulated by market forces. Contemporary researchers in the fields of organization theory and strategy concern themselves with high-level, intraorganizational entities, such as management or R&D and seem relatively untroubled by sentences in which “organization” is the subject and “learning,” the predicate. For example, Fiol and Lyles (1985) define learning, whether undertaken by individual or organizational agents, as “the process of improving actions through better knowledge and understanding.” (p. 803) Organizations learn, in the sense proposed by Leavitt and March (1988), when they “encode inferences from history into routines that guide behavior.” (p.319) Huber (1989) suggests that an organization has learned if *any of its components* have acquired information and have this information available for use, either by other components or by itself, on behalf of the organization. (p. 3)

One increasingly influential research tradition, derived from the work of Campbell (1969) and Nelson and Winter (1982), draws on the Darwinian language of evolution, adaptation, and natural selection. Researchers in this tradition see organizational learning as a process in which whole organizations or their components adapt to changing environments by generating and selectively adopting organizational routines. For example, Robert Burgelman (1994), whose work we will discuss at greater length in the following chapter, describes business firms as “ecologies of strategic initiatives.”

...which emerge in patterned ways and compete for limited organizational resources so as to increase their relative importance within the organization. Strategy results, in part, from selection and retention operating on internal variation associated with strategic initiatives. (p. 240)

In Burgeimann's description, the agents that generate and select internal variations are collective entities labeled managers, departments, or top management.

It matters greatly, of course, in Burgeimann's theory as in the theories of others mentioned above, whether entities defined at relatively high levels of social aggregation are taken to be uniquely appropriate or at least sufficient for the study of organizational adaptation and learning, or whether they are seen as needing to be complemented by a view that reveals how individuals enter into these processes. Many sociologically oriented researchers who see organizational learning as an intraorganizational phenomenon avoid the difficulties of bridging between individual and organizational phenomena by consistently treating agents and processes of learning at a relatively high level of social aggregation.

We insist, on the contrary, that a theory of organizational learning must take account of the interplay between the actions and interactions of individuals and the actions and interactions of higher-level organizational entities such as departments, divisions, or groups of managers. Unless a theory of organizational learning satisfies this criterion, it cannot contribute to knowledge useful to practitioners of organizational learning; nor can it explain the phenomena that underlie observed limitations to organizational learning. A few researchers share our view. For example, Daniel Kim (1993) has observed:

Although the meaning of the term "learning" remains essentially the same as in the individual case, the learning process is fundamentally different at the organizational level. A model of organizational learning has to resolve somehow the dilemma of imparting intelligence and learning capabilities to a nonhuman entity without anthropomorphizing it. (p. 12)

Anyone who adopts such a position faces the rather daunting task of explaining how the "fundamentally different" processes carried out by individual and by higher-level entities can interact to yield the phenomena we are prepared to recognize as organizational learning. Clearly, although organizational learning has long since become an idea in good currency, it is no less problematic for organizational theorists than when it languished at the margins of the field. Indeed, it holds a special interest for us just because it stretches the boundaries of our ordinary understandings of individual and organization.

What, then, are the possible modes of explaining the interplay between individuals and higher-level entities that constitutes organizational learning? Should we imagine that individuals who play cer-

tain organizational roles (perhaps those who exercise greatest authority or control over action) can learn from their experiences and that when enough of them do so, the organization as a whole can be said to learn? Should we think of organizations as groups of individuals, recognizing that groups are real entities irreducible to the individuals who make them up? Should we then attribute to such groups a capacity for thinking, inquiring, experimenting, and learning? Should we think of organizations as cultures that consist of systems of beliefs, values, technologies, languages, common patterns of behavior, shared representations of reality; and should we then use learning to designate certain processes of cultural change? Should we think of organizations as cognitive constructs—perhaps theories in their own right—so that organizations may be said to learn when their members contribute to the cumulative accretion or modification of these constructs?

We presented our own approach to the paradox of organizational learning in Chapter 1. A key concept for us is that of inquiry, the intertwining of thought and action carried out by individuals in interaction with one another on behalf of the organization to which they belong in ways that change the organization's theories of action and become embedded in organizational artifacts such as maps, memories, and programs. A key question for us, then, is the meaning of the phrase, "on behalf of the organization." We argue that it is possible for individuals to think and act on behalf of an organization because organizations are political entities, in a fundamental sense of that term. Collectivities become organizational when they meet three constitutional capabilities: to make collective decisions (so that groups of individuals can say "we" about themselves), to delegate authority for action to an individual in the name of the collectivity, and to say who is and who is not a member of the collectivity. Under these conditions, it makes conceptual sense to say that individuals can act on behalf of an organization. It also makes conceptual sense to say that on behalf of an organization individuals can undertake learning processes (organizational inquiry) that can, in turn, yield learning outcomes as reflected in changes in organizational theories of action and the artifacts that encode them.

But our emphasis on organizational inquiry as linking interpersonal and organizational phenomena in organizational learning and our insistence on the importance of the behavioral worlds that constrain or facilitate organizational inquiry, have led some critics of the views we first expressed in our 1978 book to dismiss our approach as one that deals exclusively with the interpersonal or social-

psychological dimension of organizational life. For some critics, indeed, our approach to organizational learning is about individuals and not about organizations at all.

Clearly, the issue underlying the controversy over our approach to organizational learning or over any attempt to treat organizational learning in terms of the interaction between individuals and organizations hinges, first, on what level or levels of aggregation one chooses to treat as distinctively organizational and, second, on the features one selects as critically important to learning at the level of aggregation in question. As we have noted, some researchers focus on clusters of organizations grouped together in larger systems, such as markets or ecologies within which learning is predicated of whole organizations (the firm or the state) or even the larger clusters to which they belong. For other observers, such as Nelson and Winter, and Burgelman, attention focuses on the interactions of larger entities (departments, divisions, top management) within organizations. For some organization theorists, such as Crozier (1963), or theorists of policy implementation (Pressman and Wildavsky, 1973; Bardach, 1974), the key focus is on games of freedoms, interests, and powers that unfold among groups of individuals who occupy kinds of roles within organizations. For other theorists, such as Hirschman, the focus is on structures of incentives, created in part through the operation of the mechanisms of exit and voice, that drive changes in organizational performance.

For theorists of a social-psychological bent, such as Schein (1992), attention focuses on individuals in interaction with one another within the settings organizations provide. Some researchers, following the directions set out in Marvin Minsky's *Society of Mind* (1991), treat individuals themselves as organizations whose thought and action must be conceived in terms of the interplay of intrapsychic microagents in direct analogy with the operation of complex computer programs.

The issue of choice of level(s) of aggregation, and the closely related issue of selective attention to features at any given level, seem to occupy in the realm of organizational phenomena a place analogous to the one they occupy in the realm of theories of material objects. Physicists, mechanical engineers, materials scientists, and physical chemists focus on strikingly different levels of aggregation (for example, galaxies, bridges, composite materials, molecules) and give privileged status to different descriptions of phenomena discovered at these levels. In part, their different foci of attention reflect what they happen to be interested in and, in part, the purposes to

which their respective inquiries are addressed. At certain key points, however, their research interests intersect, especially when the researchers are concerned with questions about the prospective guidance of technological practice. For example, civil engineers who are interested in the behavior of large-scale structures may consult materials scientists or even physical chemists, when their research leads them to think about metal fatigue or about the sources of the propagation of cracks in concrete. It remains controversial and intellectually fruitful whether descriptions of the behavior of higher-level entities, such as machines, can be reduced to descriptions of the behavior of lower-level entities, such as materials or molecules (Polanyi, 1967), or whether or in what particular ways it is both feasible and useful to develop theories of the behavior of higher-level entities without worrying much about the lower-level phenomena that might be adduced to account for that behavior.

In the broad and varied field of research on organizational learning, different interests and purposes also lead researchers to focus on different levels of aggregation and on different features of the phenomena discovered at any given level. We have argued that intersections among individual, interpersonal, and higher levels of aggregation become critically important if we wish to understand and, all the more so, if we wish to redesign the practices of organizational life, as carried out by individuals who inhabit organizations and bear responsibility for contributing to organizational performance, including especially the performance of organizational learning. But researchers on organizational learning are far from agreement on this point. There is disagreement not only about the nature of the kinds of interactions among individual, interpersonal, and higher-level entities that may be involved in organizational learning, but also whether an adequate theory of organizational learning demands an account of such interactions.

Organizational Learning Is A Meaningful Notion But Not Always Beneficent

Once organizational learning is taken as a neutral term rather than as a normative ideal, it is obvious to us, and others, that it need not be for the good, given some view of the good. In the Nazi period, Eichman's bureaucracy clearly became more efficient at carrying out its evil mission and may be said, with some plausibility, to have "learned" to do so. But the ethical critique of organizational learning varies with the kinds of evil to which the critic believes organizations are particularly disposed.

Some authors treat the ideal of the learning organization as an instance of contemporary rhetorics of "high-performance organizations," (see Kunda, 1992; who refers, in turn, to Bendix, 1956; Van Maanen, 1988; and Goffman, 1959.) They claim that organizational power elites use the ideal of the learning organization as they use other rhetorical ideals as cunning vehicles of normative control to gain the compliance, indeed, the commitment, of subordinates, and in ways that may be good for those in control but bad for those who are subordinated to them. As Kunda puts it:

Normative control is the attempt to elicit and direct the required efforts of members by controlling the underlying experiences, thoughts, and feelings that guide their actions. Under normative control, members act in the best interests of the company...[because] they are driven by internal commitment, strong identification with company goals, intrinsic satisfaction from work...elicited by a variety of managerial appeals, exhortations, and actions... In short, under normative control it is the employee's self...that is claimed in the name of corporate interest. (p. 11)

Finally, some authors criticize organizational learning because they claim that much of it, perhaps even the greater part, is in the service of stability rather than change. On this view, organizations learn to preserve the status quo, and learning of this sort is the enemy of organizational change and reform (Fiol and Lyles, 1989; Leavitt and March, 1989).

All such criticisms rest on the idea that organizational learning is not a value-neutral activity but proceeds from values, has implications for values, and is subject to critique in terms of a conception of what is good or right, and for whom. These implications, which seem obvious once they are stated, come to light only when organizational learning is stripped of its normative aura and considered as subject to evaluation in particular contexts on the basis of particular criteria of goodness or rightness. In short, we cannot escape the need to declare what kinds of organizational learning we will take to be desirable or undesirable and why.

Do Real-World Organizations Learn Productively

In order to speak of an organization learning, we must see it as a more or less coherent agent. And we must also see it as capable of acting rationally, at least in the sense of being able to remember past events, analyze alternatives, conduct experiments, and evaluate the results of action. But some authors claim that these attributions have little or no validity for organizations as we find them in the world. We

categorize their doubts in terms of threats to coherent action, valid inference, and efficacy.

Threats to Coherent Action. Some theorists have argued that organizations are actually pluralistic systems, little more than stage settings for performances by agents such as professions, disciplines, or social groupings that by their very nature cut across organizational boundaries. Some authors see organizations as political systems, made up of subgroups, each with its own interests, freedoms, and powers, crucially engaged in battles for control or avoidance of control and incapable of functioning holistically as agents of learning (Crozier, 1963; Bardach, 1974). In his middle period, March, along with various coauthors (Cohen and March, 1974; March and Olsen, 1976), proposed that organizations are inherently chaotic, at best organized anarchies. His "theory of the garbage can" presents decision making in terms of ideas, interests, images, and values in search of problems, rather than in terms of problem solvers actively searching for ideas, images, and values. Where the garbage can is in operation, it is hard to see how organizations can be considered capable of coherent action or inquiry.

Again, these lines of argument appear to have had more weight twenty years ago in the full flush of the reaction against reflective theories of organizational rationality (e.g., Perrow, 1979) than they do at present. Although attributions of organizational incoherence still present themselves as sources of doubt about claims made in the name of organizational learning, they tend no longer to be taken a priori as reasons for outright rejection of the idea. Rather, it seems, there is a growing sentiment that the degree of coherence manifested in organizational action or inquiry is an empirical matter to be ascertained at particular places and times. A case in point is March's transition from viewing organizations as "organized anarchies" to the far more modulated position he has expressed in his more recent writings, where he suggests that there are periods in which institutional reform can be pursued through "integrative processes...that treat conflict of interest as the basis for deliberation and authoritative decision rather than bargaining." (March, 1989, p. 142)

Threats to Valid Inference. Across the wide-ranging descriptions of organizational learning processes presented in scholarly literature, there is a consistent emphasis on rational inference, inference in the form of lesson drawing from observations of past experience,

inference about the causal connections between actions and outcomes, and inference from cycles of trial and error. A number of authors, including some of those noted above, base their skepticism on "threats to valid inference" which seem to them to make real-world organizational learning a dubious proposition.

March, who defines organizational learning as "encoding inferences from history into routines that guide behavior" (Leavitt and March, 1988, p. 319), has been prolific in identifying threats to the validity of such inferences. For example (1988, pp. 322-23), he undermines the importance of "competence traps," wherein organizations falsely project into the future the strategies of action that have worked for them in the past. He calls attention to various sources of ambiguity that undermine organizational judgments of success or failure:

The lessons of experience are drawn from a relatively small number of observations in a complex, changing ecology of learning organizations. What has happened is not always obvious, and the causality of events is difficult to untangle. What an organization should expect to achieve, and thus the difference between success and failure, is not always clear. (p. 323)

He describes instances of "superstitious learning" that "occur when the subjective experience of learning is compelling, but the connections between actions and outcomes are misspecified." (p. 325)

March also identifies a "dilemma of learning" that constitutes a family of threats to valid inference. When learning proceeds gradually through "small, frequent changes and inferences formed from experience with them," then a likely outcome is the reinforcement or marginal change of existing routines. Such behavior "is likely to lead to random drift rather than improvement" (Lounamaa and March, 1987). On the other hand, when organizations learn from "low probability, high consequence events," then inferences about them are often "muddled with conflict over formal responsibility, accountability, and liability" (Leavitt and March, 1989, p. 334). The upshot is that

...learning does not always lead to intelligent behavior. The same processes that yield experiential wisdom produce superstitious learning, competence traps, and erroneous inferences. (p. 335)

In this line of argument, March treats learning in the narrow sense of drawing lessons from history as an alternative to other models of decision making, such as rational choice, bargaining, and se-

lection of variations. He argues that under some circumstances learning may prove inferior to its alternatives; although he adds the caveat that the alternatives may also make mistakes, and it is, therefore, "possible to see a role for routine-based, history-dependent, target-oriented organizational learning." (p. 336) (From our point of view, all of March's alternate strategies may enter into the processes of *inquiry* around which we build our broader approach to organizational learning. The relative vulnerabilities of lesson drawing from history would be relevant, not to the general question of the cognitive capability for learning in real-world organizations, but to the problem of choosing, in any given context, what strategy of inquiry to pursue.)

A very different kind of threat to the validity of inference in organizational inquiry stems from the observation that organizational learning depends on the interpretation of events, which depends, in turn, on frames, the major story lines through which organizational inquirers set problems and make sense of experience. Framing is essential to interpretive judgments, but because frames themselves are unfalsifiable, organizational inquirers may be trapped within self-referential frames. Padgett (1992) writes that "the collectively constructed frame or 'membrane' through which information and rewards are assembled and received" is an "axiomatic construction of the world" that is "reciprocally tied to the constitution of the observer." Communication across divergent, self-referential frames is bound to be problematic.

However, Schön and Rein (1994) explore the frame conflicts that underlie persistent policy disputes, for example, those that revolve around welfare, homelessness, or the costs and benefits of advanced technology. They argue that in actual policy practice inquirers may be capable of reflective inquiry into the frames that underlie their divergent positions and can sometimes hammer out, in particular situations, a pragmatic resolution of their conflicting frames.

Threats to Effective Action. Even if organizational inquirers are sometimes able to draw valid inferences from historical experience or current observation, their inferences may not be converted to effective action. A number of contemporary researchers (Fiol and Lyles, 1985; Kim, 1993) call attention to the fact that learning outcomes may be fragmented or situational and may never enter into the organizational mainstream. In earlier research, proponents of the "behavioral theory of the firm" (Cyert and March, 1963; March, 1963; Simon, 1976) described dysfunctional patterns of organizational

behavior that undermine productive organizational learning. They noted that organizations depend on control systems which set up conflicts between rule setters and rule followers, which leads to cheating and that in such an organizational world, "everyone is rational and no one can be trusted."

"Fragmented" learning outcomes are closely related to the "conditions for error" that we described in Part II. And the dysfunctional, defensive patterns of behavior described by Simon, Cyert, and the early March are closely related to the patterns we have ascribed to limited learning systems. The question is how we should view such phenomena. Should we consider them along the lines of the behavioral theory of the firm, as pervasive and inherent features of organizational life which it is the business of organizational researchers to "discover" rather than to change? Or should we treat them as critically important impediments to productive learning that call for and may be malleable in response to, double-loop inquiry?

Conclusion

Our review of the two-pronged literature of organizational learning leaves us with challenges to the beneficence, the feasibility, and the meaningfulness of organizational learning. Proponents of the learning organization are not worried about the meaningfulness of organizational learning and take its desirability to be axiomatic. They prescribe a variety of enablers through which they claim that organizations can enhance their capability for productive learning, but they do not inquire into the gaps that separate reasonable prescription from effective implementation.

Skeptical researchers into organizational learning present from a variety of perspectives, important reasons for doubt. Some of them have raised questions about the paradox inherent in the claim that organizations learn, which hinges on assumptions about relationships among individual, interpersonal, and higher levels of social aggregation. Other writers have challenged the desirability of organizational learning, arguing that organizations may learn in ways that foster evil ends or reinforce the status quo, or arguing that the ideal of the learning organization may be used to support a subtler and darker form of managerial control. Still other researchers observe and categorize phenomena that function as impediments to valid inference and effective action.

The problems raised by the two branches of the literature are largely complementary: what one branch treats as centrally important, the other tends to ignore. Both branches do concern themselves with the capability of real-world organizations to draw valid and useful inferences from experience and observation and to convert such inferences to effective action. But authors of prescriptive bent tend to assume, uncritically, that such capabilities can be activated through the appropriate enablers, and learning skeptics tend to treat observed impediments as unalterable facts of organizational life.

In the next chapter we consider these challenges in the light of the theory-of-action perspective.