CHAPTER ONE

What is an organization that it may learn?

THE QUESTION

There has probably never been a time in our history when members, managers, and students of organizations were so united on the importance of organizational learning. Costs of health care, sanitation, police, housing, education, and welfare have risen precipitously, and we urge agencies concerned with these services to learn to increase their productivity and efficiency. Governments are torn by the conflicting demands of full employment, free collective bargaining, social welfare, and the control of inflation; we conclude that governments must learn to understand and accommodate these demands. Corporations have found themselves constrained by a web of increasingly stringent regulations for environmental protection and consumer safety, at the same time that we are most sensitive to the need for jobs and for economic growth. Government and business must learn, we say, to work together to solve these problems.

Sometimes our demands for learning turn back on our history, as when politicians and planners ask, "What have we learned from the last 20 years of housing policy?" "What have we learned from the Great Depression?" "What have we learned from Vietnam?" In a bicentennial article on "The American Experiment," Daniel Moynihan begins by asking, "What have we learned?" (Glazer and Kristol 1976)

It is not only that we are poignantly aware of our dilemmas and of the need for learning. We are also beginning to notice that there is

nothing more problematic than solutions. Some of our most agonizing problems have been triggered by our solutions to slum eradication and urban renewal, by the success of the Labor Movement in achieving income security for workers, by rising expectations consequent to our economic growth, by the unwanted consequences of technological innovations. We begin to suspect that there is no stable state awaiting us over the horizon. On the contrary, our very power to solve problems seems to multiply problems. As a result, our organizations live in economic, political, and technological environments which are predictably unstable. The requirement for organizational learning is not an occasional, sporadic phenomenon, but is continuous and endemic to our society.

Nevertheless, it is not all clear what it means for an organization to learn. Nor is it clear how we can enhance the capacity of organizations to learn.

The difficulty has first to do with the notion of learning itself. When we call for learning or change, we seem to be calling for something good. But there are kinds of change which are not good, such as deterioration, regression and stagnation. And there are kinds of learning, such as government's learning to deceive and manipulate society, which are no better. So we need to spell out both the kinds of change we have in mind when we speak of learning, and the kinds of learning we have in mind when we call for more of it.

Further, it is clear that organizational learning is not the same thing as individual learning, even when the individuals who learn are members of the organization. There are too many cases in which organizations know *less* than their members. There are even cases in which the organization cannot seem to learn what every member knows. Nor does it help to think of organizational learning as the prerogative of a man at the top who learns *for* the organization; in large and complex organizations bosses succeed one another while the organization remains very much itself, and learns or fails to learn in ways that often have little to do with the boss.

There is something paradoxical here. Organizations are not merely collections of individuals, yet there is no organization without such collections. Similarly, organizational learning is not merely individual learning, yet organizations learn only through the experience and actions of individuals.

What, then, are we to make of organizational learning? What is an organization that it may learn?

THEORY OF ACTION

In our earlier book, *Theory in Practice* (Argyris and Schon 1974), we set out to understand how practitioners of management, consultation, and intervention might learn to become more competent and effective. Our concern was especially directed to learning about interpersonal interaction. In that context, we found it useful to look at professional practice as informed by *theories of action:*

> All human beings—not only professional practitioners—need to become competent in taking action and simultaneously reflecting on this action to learn from it. The following pages provide a conceptual framework for this task by analyzing theories of action that determine all deliberate human behavior, how these theories are formed, how they come to change, and in what senses they may be considered adequate or inadequate. (p. 4)

When we attributed theories of action to human beings, we argued that all deliberate action had a cognitive basis, that it reflected norms, strategies, and assumptions or models of the world which had claims to general validity. As a consequence, human learning, we said, need not be understood in terms of the "reinforcement" or "extinction" of patterns of behavior but as the construction, testing, and restructuring of a certain kind of knowledge. Human action and human learning could be placed in the larger context of knowing.

We found it necessary to connect theories of action to other kinds of theory:

> . . . whatever else a theory of action may be, it is first a theory. Its most general properties are properties that all theories share, and the most general criteria that apply to it—such as generality, centrality, and simplicity—are criteria that apply to all theories. (p. 4)

And we also found it necessary to differentiate theories of action from theories of explanation, prediction, and control:

> A full schema for a theory of action, then, would be as follows: in situation *S*, if you want to achieve consequence *C*, under assump-

> tions *a* . . . *n*, do *A*. . . . A theory of action is a theory of deliberate human behavior which is for the agent a theory of control but which, when attributed to the agent, also serves to explain or predict his behavior. (p. 6)

Because we wished to do empirical research into human learning in situations of interpersonal interaction, we distinguished espoused theory from theory-in-use:

> When someone is asked how he would behave under certain circumstances, the answer he usually gives is his espoused theory of action for that situation. This is the theory of action to which he gives allegiance and which, upon request, he communicates to others. However, the theory that actually governs his actions is his theory-in-use, which may or may not be compatible with his espoused theory; furthermore, the individual may or may not be aware of the incompatibility of the two theories. (p. 7)

From the directly observable data of behavior, we could then ground our construction of the models of action theories which guided interpersonal behavior. And we could relate these models to the capacity for types of learning in professional practice.

It is tempting to apply this line of thought to the problem of understanding organizational learning. Perhaps organizations also have theories of action which inform their actions, espoused theories which they announce to the world and theories-in-use which may be inferred from their directly observable behavior. If so, then organizational learning might be understood as the testing and restructuring of organizational theories of action and, in the organizational context as in the individual one, we might examine the impact of models of action theories upon the capacity for kinds of learning.

But this path is full of obstacles. It is true that we do apply to organizations many of the terms we also apply to individuals. We speak of organizational action and organizational behavior. We speak also of organizational intelligence and memory. We say that organizations learn, or fail to learn. Nevertheless, a closer examination of these ways of speaking suggests that such terms are metaphors. Organizations do not literally remember, think, or learn. At least, it is not initially clear how we might go about testing whether or not they do so.

It is even puzzling to consider what it means for an organization to act or behave—notions which are essential to the construction of

organizational theories of action. Does an organization act whenever one of its members acts? If so, there would appear to be little difference between an organization and a collection of individuals. Yet it is clear that some collections of people are organizations and others are not. Furthermore, even when a collection of people is clearly an organization, individual members of the organization do many things (such as breathe, sleep, gossip with their friends) which do not seem, in some important sense, to be examples of organizational action.

If we are to speak of organizational theories of action, we must dispel some of the confusion surrounding terms like organizational intelligence, memory, and action. We must say what it means for an organization to act, and we must show how organizational action is both different from and conceptually connected to individual action. We must say what it means for an organization to know something, and we must spell out the metaphors of organizational memory, intelligence, and learning.

PERSPECTIVES ON ORGANIZATION

Let us begin by exploring several different ways of looking at an organization. An organization is:

- a government, or *polis*,
- an agency,
- a task system.

Each of these perspectives will illuminate the sense in which an organization may be said to act. Further, an organization is:

- a theory of action,
- a cognitive enterprise undertaken by individual members,
- a cognitive artifact made up of individual images and public maps.

Each of these descriptions will reveal the sense in which an organization may be said to know something, and to learn.

Consider a mob of students protesting against their university's policy. At what point do they cease to be a mob and begin to be an organization?

The mob is a collectivity. It is a collection of people who may run, shout, and mill about together. But it is a collectivity which cannot make a decision or take an action in its own name, and its boundaries are vague and diffuse.

As the mob begins to meet three sorts of conditions, it becomes more nearly an organization. Members must devise procedures for: (1) making decisions in the name of the collectivity, (2) delegating to individuals the authority to act for the collectivity, and (3) setting boundaries between the collectivity and the rest of the world. As these conditions are met, members of the collectivity begin to be able to say "we" about themselves; they can say, "We have decided," "We have made our position clear," "We have limited our membership." There is now an organizational "we" that can decide and act.

When the members of the mob become an identifiable vehicle for collective decision and action, they become, in the ancient Greek sense of the term, a *polis*, a political entity. Before an organization can be anything else, it must be in this sense political, because it is as a political entity that the collectivity can take organizational action. It is individuals who decide and act, but they do these things *for* the collectivity by virtue of the rules for decision, delegation, and membership. When the members of the collectivity have created such rules, they have organized.

Rule making need not be a conscious, formal process. What is important is that members' behavior be rule-governed in the crucial respects. The rules themselves may remain tacit, unless for some reason they are called into question. So long as there is continuity in the rules which govern the behavior of individuals, the organization will persist, even though members come and go. And what is most important for our purposes, it now becomes possible to set up criteria of relevance for constructing organizational theory-in-use. Organizational theory-in-use is to be inferred from observation of organizational behavior—that is, from organizational decisions and actions. The decisions and actions carried out by individuals are organizational insofar as they are governed by collective rules for decision and delegation. These alone are the decisions and actions taken in the name of the organization.

Through such a process, a mob becomes an organization. But if we are interested in organizational theory of action, we must ask what *kind* of an organization it becomes.

If a collection of people begins to decide and to act on a continuing basis, it becomes an instrument for continuing collective action, an *agency*. In this sense, the collections of workers involved in the labor movement organized from time to time to form unions, and collections of individual investors organized to form limited liability corporations. Such agencies have functions to fulfill, work to do. Their theories-in-use may be inferred from the ways in which they go about doing their work.

Generally speaking, an agency's work is a complex task, continually performed. The agency—an industrial corporation, a labor union, a government bureau, or even a household—embodies a strategy for decomposing that complex task into simpler components which are regularly delegated to individuals. Organizational roles—president, lathe-operator, shop steward—are the names given to the clusters of component tasks which the agency has decided to delegate to individual members. The organization's *task system*, its pattern of interconnected roles, is at once a design for work and a division of labor.

An agency is thus the solution to a problem. It is a strategy for performing a complex task which might have been carried out in other ways. This is true not only for the design of the task system, the division of labor, but also for the selection of strategies for performing component tasks.

We can view a sugar-refining company, for example, as an answer to questions such as these: What is the best way to grow and harvest cane? How should it be refined? How is it best distributed and marketed? For each subquestion the organization is an answer. The company's way of growing cane reflects certain strategies (for the cultivation of land, for harvesting and fertilizing), certain norms (for productivity and quality, for the use of labor), and certain assumptions (about the yields to be expected from various patterns of cultivation). The norms, strategies, and assumptions embedded in the company's cane-growing practices constitute its *theory of action* for cane-growing. There are comparable theories of action implicit in the company's ways of distributing and marketing its products. Taken together, these component theories of action represent a theory of action for achieving corporate objectives. This global theory of action we call "instrumental." It includes norms for corporate performance (for example, norms for margin of profit and for return on invest-

ment), strategies for achieving norms (for example, strategies for plant location and for process technology), and assumptions which bind strategies and norms together (for example, the assumption that maintenance of a high rate of return on investment depends on the continual introduction of new technologies).

The company's instrumental theory of action is a complex system of norms, strategies, and assumptions. It includes in its scope the organization's patterns of communication and control, its ways of allocating resources to goals, and its provisions for self-maintenance—that is, for rewarding and punishing individual performance, for constructing career ladders and regulating the rate at which individuals climb them, and for recruiting new members and instructing them in the ways of the organization.

Like the rules for collective decision and action, organizational theories of action need not be explicit. Indeed, formal corporate documents such as organization charts, policy statements, and job descriptions often reflect a theory of action (the *espoused theory*) which conflicts with the organization's *theory-in-use* (the theory of action constructed from observation of actual behavior)—and the theory-in-use is often tacit. Organizational theory-in-use may remain tacit, as we will see later on, because its incongruity with espoused theory is *undiscussable*. Or it may remain tacit because individual members of the organization know more than they can say—because the theory-in-use is *inaccessible* to them. Whatever the reason for tacitness, the largely tacit theory-in-use accounts for organizational identity and continuity.

Consider a large, enduring organization such as the U.S. Army. Over 50 years or so, its personnel may turn over completely, yet we still speak of it as "the Army." It is no longer the same collection of people, so in what sense is it still the same? Suppose we wanted to discover whether it was in fact the same organization. How would we proceed? We might examine uniforms and weapons, but in 50 years these might have changed entirely. We might then study the 50-year evolution of military practices—that is, the norms for military behavior, the strategies for military action, the assumptions about military functioning. We would then be studying the evolution of the Army's theory-in-use. And we might learn that certain features of it—for example, the pattern of command, the methods of training, the division into regiments and platoons—had remained essentially

unchanged, while other features of it—battle strategies, norms for performance—had evolved continuously from earlier forms. We might conclude that we were dealing with a single organization, self-identical, whose theory-in-use had evolved considerably over time.

It is this theory-in-use, an apparently abstract thing, which is most distinctively real about the Army. It is what old soldiers know and new ones learn through a continuing process of socialization. And it is the history of change in theory-in-use which we would need to consult in order to inquire into the Army's organizational learning.

In order to discover an organization's theory-in-use, we must examine its practice, that is, the continuing performance of its task system as exhibited in the rule-governed behavior of its members. This is, however, an outside view. When members carry out the practices appropriate to their organization, they are also manifesting a kind of knowledge. And this knowledge represents the organization's theory-in-use as seen from the inside.

IMAGES AND MAPS

Each member of the organization constructs his or her own representation, or image, of the theory-in-use of the whole. That picture is always incomplete. The organization members strive continually to complete it, and to understand themselves in the context of the organization. They try to describe themselves and their own performance insofar as they interact with others. As conditions change, they test and modify that description. Moreover, others are continually engaged in similar inquiry. It is this continual, concerted meshing of individual images of self and others, of one's own activity in the context of collective interaction, which constitutes an organization's knowledge of its theory-in-use.

An organization is like an organism each of whose cells contains a particular, partial, changing image of itself in relation to the whole. And like such an organism, the organization's practice stems from those very images. Organization is an artifact of individual ways of representing organization.

Hence, our inquiry into organizational learning must concern itself not with static entities called organizations, but with an active process of organizing which is, at root, a cognitive enterprise. Individual members are continually engaged in attempting to know the

organization, and to know themselves in the context of the organization. At the same time, their continuing efforts to know and to test their knowledge represent the object of their inquiry. Organizing is reflexive inquiry.

From this perspective, organizational continuity is a considerable achievement. But we could not account for organizational continuity if the cognitive enterprise of organizing were limited to the private inquiry of individuals. Even when individuals are in face-to-face contact, private images of organization erode and diverge from one another. When the task system is large and complex, most members are unable to use face-to-face contact in order to compare and adjust their several images of organizational theory-in-use. They require external references. There must be public representations of organizational theory-in-use to which individuals can refer.

This is the function of organizational maps. These are the shared descriptions of organization which individuals jointly construct and use to guide their own inquiry. They include, for example, diagrams of work flow, compensation charts, statements of procedure, even the schematic drawings of office space. A building itself may function as a kind of map, revealing patterns of communication and control. Whatever their form, maps have a dual function. They describe actual patterns of activity, and they are guides to future action. As musicians perform their scores, members of an organization perform their maps.

Organizational theory-in-use, continually constructed through individual inquiry, is encoded in private images and in public maps. These are the media of organizational learning.

ORGANIZATIONAL LEARNING

As individual members continually modify their maps and images of the organization, they also bring about changes in organizational theory-in-use.

Not all of these changes qualify as learning. Members may lose enthusiasm, become sloppy in task performance, or lose touch with one another. They may leave the organization, carrying with them important information which becomes lost to the organization. Or changes in the organization's environment (a slackening of demand for product, for example) may trigger new patterns of response which

undermine organizational norms. These are kinds of deterioration, sometimes called organizational entropy.

But individual members frequently serve as agents of changes in organizational theory-in-use which run counter to organizational entropy. They act on their images and on their shared maps with expectations of patterned outcomes, which their subsequent experience confirms or disconfirms. When there is a mismatch of outcome to expectation (error), members may respond by modifying their images, maps, and activities so as to bring expectations and outcomes back into line. They detect an error in organizational theory-in-use, and they correct it. This fundamental learning loop is one in which individuals act from organizational theory-in-use, which leads to match or mismatch of expectations with outcome, and thence to confirmation of disconfirmation of organizational theory-in-use.

Quality control inspectors detect a defect in product, for example; they feed that information back to production engineers, who then change production specifications to correct that defect. Marketing managers observe that monthly sales have fallen below expectations; they inquire into the shortfall, seeking an interpretation which they can use to devise new marketing strategies which will bring the sales curve back on target. When organizational turnover of personnel increases to the point where it threatens the steady performance of the task system, managers may respond by investigating the sources of worker dissatisfaction; they look for factors they can influence—salary levels, fringe benefits, job design—so as to reestablish the stability of their work force.

Single-loop learning

In these examples, *members of the organization respond to changes in the internal and external environments of the organization by detecting errors which they then correct so as to maintain the central features of organizational theory-in-use.* These are learning episodes which function to preserve a certain kind of constancy. As Gregory Bateson has pointed out (Bateson 1972), the organization's ability to remain stable in a changing context denotes a kind of learning. Following his usage, we call this learning single-loop. (Bateson 1960) There is a single feed-back loop which connects detected outcomes of action to organizational strategies and assumptions which are modi-

fied so as to keep organizational performance within the range set by organizational norms. The norms themselves—for product quality, sales, or task performance—remain unchanged.

These examples also help to make clear the relationship between individual and organizational learning. The key to this distinction is the notion of *agency. Just as individuals are the agents of organizational action, so they are the agents for organizational learning.* Organizational learning occurs when individuals, acting from their images and maps, detect a match or mismatch of outcome to expectation which confirms or disconfirms organizational theory-in-use. In the case of disconfirmation, individuals move from error detection to error correction. Error correction takes the form of inquiry. The learning agents must discover the sources of error—that is, they must attribute error to strategies and assumptions in existing theory-in-use. They must invent new strategies, based on new assumptions, in order to correct error. They must produce those strategies. And they must evaluate and generalize the results of that new action. "Error correction" is shorthand for a complex learning cycle.

But in order for *organizational* learning to occur, learning agents' discoveries, inventions, and evaluations must be embedded in organizational memory. They must be encoded in the individual images and the shared maps of organizational theory-in-use from which individual members will subsequently act. If this encoding does not occur, individuals will have learned but the organization will not have done so.

Suppose, for example, that the quality control inspectors find a product defect which they then decide to keep to themselves, perhaps because they are afraid to make the information public. Or suppose that they try to communicate this information to the production engineers, but the production engineers do not wish to listen to them. Or suppose that the interpretation of error requires collaborative inquiry on the part of several different members of the organization who are unwilling or unable to carry out such a collaboration. (Indeed, because organizations are strategies for decomposing complex tasks into task/role systems, error correction normally requires collaborative inquiry.) In all of these instances, individual learning may or may not have occurred, but individuals do not function as agents of organizational learning. What individuals may have learned remains as an unrealized potential for organizational learning.

From this it follows both that there is no organizational learning without individual learning, and that individual learning is a necessary but insufficient condition for organizational learning. We can think of organizational learning as a process mediated by the collaborative inquiry of individual members. In their capacity as agents of organizational learning, individuals restructure the continually changing artifact called organizational theory-in-use. Their work as learning agents is unfinished until the results of their inquiry—their discoveries, inventions, and evaluations—are recorded in the media of organizational memory, the images and maps which encode organizational theory-in-use.

If we should wish to test whether organizational learning has occurred, we must ask questions such as these: Did individuals detect an outcome which matched or mismatched the expectations derived from their images and maps of organizational theory-in-use? Did they carry out an inquiry which yielded discoveries, inventions, and evaluations pertaining to organizational strategies and assumptions? Did these results become embodied in the images and maps employed for purposes such as control, decision, and instruction? Did members subsequently act from these images and maps so as to carry out new organizational practices? Were these changes in images, maps, and organizational practices regularized so that they were unaffected by some individual's departure? Do new members learn these new features of organizational theory of action as part of their socialization to the organization?

Each of these questions points to a possible source of failure in organizational learning, as well as to the sources of organizational learning capacity. So far, however, we have limited ourselves to the kind of learning called single-loop. Let us now consider learning of another kind.

Double-loop learning

Organizations are continually engaged in transactions with their internal and external environments. Industrial corporations, for example, continually respond to the changing pattern of external competition, regulation and demand, and to the changing internal environment of workers' attitudes and aspirations. These responses take the form of error detection and error correction. Single-loop learning is sufficient where error correction can proceed by changing

organizational strategies and assumptions within a constant framework of norms for performance. It is concerned primarily with effectiveness—that is, with how best to achieve existing goals and objectives and how best to keep organizational performance within the range specified by existing norms. In some cases, however, error correction requires an organizational learning cycle in which organizational norms themselves are modified.

Consider an industrial firm which has set up a research and development division charged with the discovery and development of new technologies. This has been a 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. But the new division generates technologies which do not fit the corporation's familiar pattern of operations. In order to exploit some of these technologies, for example, the corporation may have to turn from the production of intermediate materials with which it is familiar to the manufacture and distribution of consumer products with which it is unfamiliar. But 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 the business they are in. 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 operations constant, they must give up on the imperative for growth, at least 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. They must, to begin with, recognize the conflict itself. They have set up a new division which has yielded unexpected outcomes; this is an error, in the sense

earlier described. They must reflect upon this error to the point where they become aware that they cannot correct it by doing better what they already know how to do. They must become aware, for example, that they cannot correct the error by getting the new division to perform more effectively under existing norms; indeed, the more effective the new division is, the more its results will plunge the managers into 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 which 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, which must then be embedded in the images and maps which encode organizational theory-in-use.

We call this sort of learning *double-loop*. There is in this sort of episode a double feedback loop with connects the detection of error not only to strategies and assumptions for effective performance but to the very norms which define effective performance.

Single-loop learning, as we have defined it, consists not only of a change in organizational strategies and assumptions but of the particular sort of change appropriately described as learning. In single-loop learning, members of the organization carry out a collaborative inquiry through which they discover sources of error, invent new strategies designed to correct error, produce those strategies, and evaluate and generalize the results. Similarly, double-loop learning consists not only of a change in organizational norms but of the particular sort of inquiry into norms which is appropriately described as learning.

In organizational double-loop learning, incompatible requirements in organizational theory-in-use are characteristically expressed through a conflict among members and groups within the organization. In the industrial organization, for example, some managers may become partisans of growth through research and of a new image of the business based upon research, while others may become opponents of research through their allegiance to familiar and predictable patterns of corporate operation. Double-loop learning, if it occurs,

will consist of the process of inquiry by which these groups of managers confront and resolve their conflict.

In this sense, the organization is a medium for translating incompatible requirements into interpersonal and intergroup conflict.

Members of the organization may respond to such a conflict 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 are to be made among competing requirements, and weightings and priorities are to be set on the basis of prevailing power. The "R & D faction," for example, may include the chief executive who wins out over the "old guard" through being more powerful. Or the two factions may fight it out to a draw, settling their differences in the end by a compromise which 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. The conflict is settled not by inquiry but by fighting it out. Neither side emerges from the settlement with a new sense of the nature of the conflict, of its causes and consequences, or of its meaning for organizational theory-in-use.

On the other hand, parties to the conflict may engage the conflict through inquiry of the following kinds:

- They may invent new strategies of performance which circumvent the perceived incompatibility of requirements. They may, for example, succeed in defining a kind of research and development addressed solely to the existing pattern of business, which offers the likelihood of achieving existing norms for growth. They will then have succeeded in finding a single-loop solution to what at first appeared as a double-loop problem.
- They may carry out a "trade-off analysis" which enables them to conclude jointly that so many units of achievement of one norm are balanced by so many units of achievement of another. 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 limited 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 the options for action.

- In the context of the conflict, the incompatible requirements may not lend themselves to trade-off analysis. They may be perceived as incommensurable. In such a case, the conflict may still be resolved through inquiry which gets underneath the members' starting perceptions of the incompatible requirements. Participants must then ask why they hold the positions they do, and what they mean by them. They may ask, what factors have led them to adopt these particular standards for growth in sales and earnings, what their rationale is, and what are likely to be the consequences of attempting to achieve them, through any means whatever? Similarly, they may ask what kinds of predictability in operations are of greatest importance, to whom they are most important, and what conditions make them important.

Such inquiry may lead to a significant restructuring of the configuration of corporate norms. Or it may lead to the invention of new patterns of incentives, budgeting, and control which take greater account of requirements for both growth and predictability.

We will give the name "double-loop learning" to those sorts of organizational inquiry which resolve imcompatible organizational norms by setting new priorities and weightings of norms, or by restructuring the norms themselves together with associated strategies and assumptions.

In these cases, individual members resolve the interpersonal and intergroup conflicts which express incompatible requirements by creating new understandings of the conflicting requirements, their sources, conditions, and consequences—understandings which then become embedded in the images and maps of organization. By doing so, they make the new, more nearly compatible requirements susceptible to effective realization.

There are three observations we wish to make about distinction between single- and double-loop organizational learning. They will become clearer through our discussion of examples in later chapters but need brief mention here.

First, it is often impossible, in the real-world context of organizational life, to find inquiry cleanly separated from the uses of power.

Inquiry and power-play are often combined. Some of the ways in which they are combined, and some of the ways in which power-play inhibits inquiry, we will consider at length in Part II. Given such mixtures, we will want to differentiate the two kinds of processes which are often mixed in practice so that we may speak of those aspects of interpersonal and intergroup conflict which involve organizational learning and those which do not.

Second, while we have described the *kinds* of inquiry which are essential to single- and double-loop learning, we have not yet dwelt on the *quality* of inquiry. Two different examples of double-loop learning, both of which exhibit detection of error and correction of error through the restructuring of organizational norms, may be of unequal quality. The same is true of single-loop learning. Organizations may learn more or less well, yet their inquiries may still qualify as learning of the single- or double-loop kind. The standards we mean to apply to such judgments of quality will occupy us in later chapters, particularly in Chapter 6.

Finally, we must point out that the distinction between single- and double-loop learning is less a binary one than might first appear. Organizational theories-in-use are systemic structures composed of many interconnected parts. We can examine these structures from the point of view of a particular, local theory of action, such as the industrial firm's theory of action for quality control, or we can attend to more global aspects of the structure, such as the firm's theory of action for achieving targeted return on investment. Furthermore, certain elements are more fundamental to the structure and others are more peripheral. For example, the industrial firm's norms for growth and for predictability of management are fundamental to its theory of action—in the sense that if they were changed, a great deal of the rest of the theory of action would also have to change—and it is their fundamental status which gives a special poignancy to their conflict. On the other hand, a particular norm for product quality may be quite peripheral to the organization's theory of action; it could change without affecting much of the rest of the theory of action.

Now, an inquiry into a *strategy* fundamental to the firm's theory of action, such as the strategy of measuring divisional performance by monthly profit-and-loss statements, will be likely to involve much of the rest of the organization's theory of action, including its norms. But an inquiry into a *norm* peripheral to the organization's theory of ac-

tion may involve very little of the rest of its theory of action. From this, two conclusions may be drawn. First, in judging whether learning is single- or double-loop, it is important to notice where inquiry goes as well as where it begins. Second, it is possible to speak of organizational learning as *more or less* double-loop. In place of the binary distinction we have a more continuous concept of depth of learning.

It is possible, we think, to make clear distinctions between relatively deep and relatively peripheral examples of organizational learning. We will continue to call the former double- and the latter, single-loop learning. Our examples of double-loop learning will involve norms fundamental to organizational theories of action, for these are the examples we believe to be of greatest importance. The reader should keep in mind, however, that we speak of these categories as discrete when they are actually parts of a continuum.

With these *caveats*, we can return to our main line of argument.

Deutero-learning

Since World War II, it has gradually become apparent not only to business firms but to all sorts of organizations that the requirements of organizational learning, especially for double-loop learning, are not one-shot but continuing. There has been a sequence of ideas in good currency—such as "creativity," "innovation," "the management of change"—which reflect this awareness.

In our earlier example, to take one instance, managers of the industrial firm might conclude that their organization needs to learn how to restructure itself, at regular intervals, so as to exploit the new technologies generated by research and development. That is, the organization needs to learn how to carry out single- and double-loop learning.

This sort of learning to learn Gregory Bateson has called *deutero-learning* (that is, second-order learning). Bateson illustrates the idea through the following story:

> A female porpoise . . . is trained to accept the sound of the trainer's whistle as a "secondary reinforcement." The whistle is expectably followed by food, and if she later repeats what she was doing when the whistle blew, she will expect again to hear the whistle and receive food.
>
> The porpoise is now used by the trainers to demonstrate "operant conditioning" to the public. When she enters the exihibition tank,

> she raises her head above the surface, hears the whistle and is fed . . .
>
> But this pattern is (suitable) only for a single episode in the exhibition tank. She must break that pattern to deal with the *class* of such episodes. There is a larger context of contexts which will put her in the wrong . . .
>
> When the porpoise comes on stage, she again raises her head. But she gets no whistle. The trainer waits for the next piece of conspicuous behavior, likely a tail flip, which is a common expression of annoyance. This behavior is then reinforced and repeated (by giving her food).
>
> But the tail flip was, of course, not rewarded in the third performance.
>
> Finally the porpoise learned to deal with the context of contexts—by offering a different or new piece of conspicuous behavior whenever she came on stage.

Each time the porpoise learns to deal with a larger class of episodes, she learns *about* the previous contexts for learning. Her creativity reflects deutero-learning.

When an organization engages in deutero-learning, its members learn, too, about previous contexts for learning. They reflect on and inquire into previous contexts for learning. They reflect on and inquire into previous episodes of organizational learning, or failure to learn. They discover what they did that facilitated or inhibited learning, they invent new strategies for learning, they produce these strategies, and they evaluate and generalize what they have produced. The results become encoded in individual images and maps and are reflected in organizational learning practice.

The deutero-learning cycle is relatively familiar in the context of organizational learning curves. Aircraft manufacturers, for example, project the rate at which their organizations will learn to manufacture a new aircraft and base cost estimates on their projections of the rate of organizational learning. In the late 1950s, the Systems Development Corporation undertook the "Cogwheel" experiment, in which members of an aircraft-spotting team were invited to inquire into their own organizational learning and then to produce conditions which would enable them more effectively to learn to improve their performance (Chapman and Kennedy 1956).

In these examples, however, deutero-learning concentrates on single-loop learning; emphasis is on learning for effectiveness rather than on learning to resolve conflicting norms for performance. But the concept of deutero-learning is also relevant to double-loop learning. How, indeed, can organizations learn to become better at double-loop learning? How can members of an organization learn to carry out the kinds of inquiry essential to double-loop learning? What are the conditions which enable members to meet the tests of organizational learning? And how can they learn to produce those conditions?

Organizations are not only theories of action. They are also small societies composed of persons who occupy roles in the task system. What we have called the internal environment of an organization is the society of persons who make up the organization at any given time. These societies have their own characteristic behavioral worlds. These enable us to recognize a person as "an army man," "a government man," "a General Electric man." Within these societies, members tend to share characteristic languages, styles, and models of *individual* theory-in-use for interaction with others. In the light of these behavioral worlds, we can and do describe organizations as more or less "open," "experimental," "confronting," "demanding," or "defensive." These behavioral worlds, with their characteristic models of individual theory-in-use, may be more or less conducive to the kinds of collaborative inquiry required for organizational learning.

Hence, if we wish to learn more about the conditions that facilitate or inhibit organizational learning, we must explore the ways in which the behavioral worlds of organizations affect the capacity for inquiry into organizational theory-in-use.

SUMMARY

Organizational learning is a metaphor whose spelling out requires us to reexamine the very idea of organization. A collection of individuals organizes when its members develop rules for collective decision delegation and membership. In their rule-governed behavior, they act for the collectivity in ways that reflect a task system. Just as individual theories of action may be inferred from individual behavior, so organizational theories of action may be inferred from patterns of organizational action. As individuals have espoused theories which may

be incongruent with their (often tacit) theories-in-use, so with organizations.

Organizational learning occurs when members of the organization act as learning agents for the organization, responding to changes in the internal and external environments of the organization by detecting and correcting errors in organizational theory-in-use, and embedding the results of their inquiry in private images and shared maps of organization.

In organizational single-loop learning, the criterion for success is effectiveness. Individuals respond to error by modifying strategies and assumptions within constant organizational norms. In double-loop learning, response to detected error takes the form of joint inquiry into organizational norms themselves, so as to resolve their inconsistency and make the new norms more effectively realizable. In both cases, organizational learning consists of restructuring organizational theory of action.

When an organization engages in deutero-learning, its members learn about organizational learning and encode their results in images and maps. The quest for organizational learning capacity must take the form of deutero-learning; most particularly about the interactions between the organization's behavioral world and its ability to learn.

CHAPTER TWO

What facilitates or inhibits organizational learning?

THE MERCURY STORY

Let us consider the case of a corporation which we will call the Mercury Corporation.

It is a chemical company, initially formed in the 1920s around a single product line. It was among the first in its industry to set up a research and development division, and over the decades it has generated new businesses in many different fields. By the mid-1960s it had grown to over five billion dollars in sales. Management, preoccupied with maintaining the company's rate of growth, began to worry about signs of flagging vitality. They localized the source of the difficulty in the declining capacity of the research and development division to generate new products which could become the basis for new businesses.

They diagnosed the problem as "an entrepreneurial gap." Research, they thought, had the capacity to generate new technologies, but not to commercialize them. For this, Research depended on existing divisions of the company which were often reluctant to risk a very new product or process.

Given this diagnosis of the problem, management invented a solution that seemed appropriate. They would establish a New-Business Division which would be empowered not only to develop new technologies but to "incubate" them. The new division would be able

to make and sell new products, turning them over to existing divisions only when they had already proved their worth. Existing divisions, freed from the need to shoulder the risks of the first reduction to practice, would no longer resist the new technologies generated by Research.

In the early 1970s Mercury experienced a different kind of crisis. The phenomenal growth of the corporation seemed now to go hand in hand with a declining rate of earnings. Management attributed this to the new scale of the business. The corporation's success in achieving its targets for growth had created an enterprise whose size and scope exceeded the capacity of the central administration. A business in excess of five billion dollars could not be run in the same old way. There was a decision to decentralize the management of the corporation, creating semiautonomous divisions, each with its own business charter and its own president. These presidents became known as the "barons."

By the mid-1970s the New-Business Division had accumulated 10 years of experience. It had accounted for some 20 million dollars in research and development expenditure without giving rise to a new business of any consequence. In spite of this record, the new Division remained in being. Indeed, top management had not directly criticized the Division. But members of the NBS staff were becoming uneasy, and they called in an outside consultant.

The consultant set out to map the problem. He found that various members of the organization held different and sometimes conflicting views of the problem, for example:

> "It's a corporate weakness; we know how to manage capital better than we do technology."
>
> "We're constrained by our charter; top management needs to free us up so that, if necessary, we can cut across the charters of existing divisions."
>
> "We've taken the path of least resistance. We need to be more 'gutsy,' to take the business risks we were set up to take."
>
> "We need to do much more of what we have been doing, only with more creativity. We need lots of balls in the air."

Individuals expressed such views in private but not in public. They had never confronted one another with their differences, nor

had they pursued their disagreements to the point of asking what might test them.

The consultant proposed to get underneath these differences by asking his clients—the Vice President for Technology, the Director of Research, members of the NBD staff—to construct case histories of corporate development over the past decade, both the developments they saw as successes and those they saw as failures.

This turned out not to be an easy task. For one thing, the failures tended to remain buried; people were not eager to exhume them. For another, individuals who held parts of the stories in their own heads had never concerted or sought to reconcile their differing pictures of events. Knowledge of past efforts at development remained scattered among members of the corporation.

In consequence, the corporation had been able to create certain myths of development. These included, for example, the notions that "Good ideas come from the top," and "A good idea will find its own way." These myths were widely diffused throughout the organization, and they remained untested. Indeed, they became self-fulfilling prophecies. Successes were read, after the fact, as good ideas, and failures were read as bad ones. Because failures remained buried, the quality of their ideas was never compared with the quality of the successes. Moreover, success itself came to be accepted as the test of a good idea. Because (as we will see) top management support was essential to success, it was easy to believe that good ideas came from the top.

Nevertheless, with the consultant's urging and support, members of the corporation came together to construct case histories of new business development. The following, in brief, is the story of a "success":

> In the Resin Division, a new manufacturing process had been developed which the Resin Division itself could not use. But the process looked to the research director as though it might have potential for textile products, and there was a small Textiles Division. A task force was set up to conduct a feasibility study.
>
> At first, it was not evident to the Resin Division that they might *lose* market through the new development, as they lost old customers with whom the corporation would now be competing. So they contributed staff to the task force, as did the Textiles and Research Divisions. This was a drain on divisional resources, and

it hurt. But the chairman of the board was behind the effort. He made corporate funds available, and he smoothed the way. In addition, about midway through the study a big customer came along and said, "If you can make it, I'll buy it."

In the course of the construction of the story, the following dialogue occurred among the consultant, a vice-president, and a division manager:

C: Why hadn't they told the head of the Resin Division that he might actually lose market?

VP: If we'd told him, given our management policies, he would have been a fool to go along.

C: Did you test that assumption?

VP: No. You can't test something like that.

C: What happened when he found out?

VP: He blew his top, but he calmed down.

DM: And some of us wondered when someone was going to play that game on us.

C: Did this make you wary and mistrustful?

DM: You're right. You have to be that way if you're going to survive.

Other "successes" fell into a similar pattern. First, there was the perception of a business opportunity based on the development of a new technology. The technology usually grew out of an existing division; often, it was the unexpected consequence of planned R & D activity. The exploitation of the business opportunity required the recombination of existing businesses. Members of the corporation noticed, for example, that: "If we could only group these businesses together, we have a family of technologies that could help them all to take off," or "Here's a technology that would be good for fastening. We don't see ourselves now in the fastening business, but there are pieces of other divisions which have to do with fasteners. If we combined them in a new division, we'd have the marketing vehicle for our new technology."

Often, existing divisions resisted giving up a piece of themselves for the sake of creating a new business for the corporation. As one member said, "In every successful marriage, there's at least one

unwilling partner!" As long as possible, promoters of the new business met that resistance by working around it and by withholding information. When those strategies were no longer feasible, there was recourse to top management. Top management would then make commitments to the new business, make corporate funds available for development, and talk one of the barons into going along.

Over the previous decade, three such successes had occurred, and they accounted for a major share of corporate growth.

"Failures" also fell into a pattern.

In these stories, too, technological opportunities had been detected. But here the technology was seen as "a way of entering a business we're not yet in." The new business had then to be described to top management. Typically, top management gave one of two responses. Sometimes they saw the new business as "too small" compared to the standard of size set by existing businesses; promoters of the new business found it difficult to make the case that a new and as yet untried business would meet those criteria. On the other hand, top management might see the new business as "too unfamiliar," "not our kind of thing." Even if the business seemed to promise a large enough volume of sales, the process of getting from here to there seemed too uncertain. When top management felt uncomfortable, they refused to make any commitment to the project.

Without top management commitment, even though some resources might be devoted to development, there was not enough energy to bring the project to fruition. The development would be dropped. Failure might then be attributed to "a bad idea" or to someone's "incompetence." As a consequence, the best technical people wanted to stay away from such developments, and those who had to stay with them acquired what someone called a "loser syndrome."

Taking the two kinds of story together, the consultant and the client group put together the following picture.

The imperative for increasing R & D productivity had led to the creation of the New-Business Division. The great increase in corporate size had led to decentralization of management. The new scale of corporate sales also set massive requirements for the volume of sales that would have to be delivered by a promising new business.

The divisions, with their semiautonomous status, operated under strict incentives to deliver profit. They held fiercely to their business charters, even when they were not fully exploiting them. They were

reluctant to give up a piece of themselves for the sake of creating a new business for the corporation.

But given the standards for new businesses set by corporate scale, and the requirement for familiarity that made top management comfortable with a new business, promising new businesses were seen only in the recombination of pieces of the existing business. Only these were familiar enough to make top management comfortable, and big enough to make a still bigger business believable. Hence it was only the recombinations that could get enough support from top management to make the barons see themselves as pieces of the corporate puzzle.

Hence the pattern of the "successes."

But each such success left the barons more mistrustful of central administration, more wary of getting caught up in such losing games.

The New-Business Division was based on the theory that research would yield new technologies which had then to be combined with business opportunities. These "balls in the air" could then be screened according to corporate criteria, and the most promising of them selected for further investment.

However, NBD existed as an island surrounded by hostile divisional barons, every one of whom saw NBD as a threat to divisional integrity and a drain on divisional resources. NBD sought to avoid confrontation with the barons. They sought to protect themselves by going after new businesses which lay *outside* of divisional business charters. Hence, they spent their time exploring opportunities which lay outside the existing business. But NBD depended ultimately on top management for the support to carry a project forward, and top management saw these opportunities as either too small or too unfamiliar.

Hence the pattern of the "failures."

Moreover, the whole process remained undiscussable. For if it were to be discussed, it would have been necessary to bring out into the open the games played between central administration and the divisions, and the barons' competition with NBD. It would also have been necessary to face publicly the conflicting requirements of corporate development and of decentralized management. The first meant that each divisional manager must see himself, at least potentially, as a piece of the corporate puzzle, ready to give up resources and markets for the sake of a new corporate business. The second meant

that each divisional manager must exploit and protect his business territory, and respond only to the demands, month by month and year by year, that he deliver expected earnings.

Hence the absence of direct criticism of NBD. No one was anxious to open that can of worms.

PERSPECTIVES ON ORGANIZATIONAL LEARNING

The Mercury story is both a story of organizational learning, and of organizational nonlearning. It illustrates concepts we have already introduced and others not yet defined which are important to the inquiry we intend to pursue.

At the heart of the story is an inconsistency in Mercury's corporate theory of action, an inconsistency generated by the company's very success in achieving corporate growth. In the early 1970s, that success had produced a situation which management had then defined as a problem of organizational structure: central administration overloaded by the task of managing the business. Management also found a structural solution to that problem: decentralization. And its theory-in-use for decentralization included the norm for central retention of control over corporate performance, even while divisional managers were being freed up to manage their own divisional businesses. The strategy of control was to set targets for divisional performance as measured by earnings, and to use the full array of corporate incentives and sanctions to reward the meeting of targets and to punish failure. Divisional managers' theory-in-use for responding to central control included self-protection by exercising their own unilateral control wherever they could do so, and by defending their charters and territories against all comers.

In the middle 1960s, however, management had constructed a different sort of problem: that of a decline in corporate rate of growth. In this case, the pattern of solution had taken more than one form. On the basis of its diagnosis of the "entrepreneurial gap," top management had again created a structural solution: the New-Business Division. This structural invention was based on an espoused theory for corporate development. New business opportunities were to be generated and screened; the most promising were to be incubated by New-Business Division, then passed on to existing divisions for management. To avoid conflict with existing divisions, the new Divi-

sion developed a theory-in-use for its activities which included the strategy of seeking new opportunities outside the bounds of existing business. But at the same time the corporation manifested another kind of theory-in-use for the development of new businesses, one which employed the strategy of seeking new business opportunities from the recombination of pieces of the existing business, well within the zone that top management would find both credible and comfortable. But this theory-in-use carried with it the requirement that divisional managers, from time to time, give up parts of themselves to new corporate developments.

Thus the inconsistency in the requirements for decentralized management and for corporate development came to a head around the interaction of divisional with central management. Divisional managers were, as they saw it, subject to inconsistent corporate demands. The inconsistency in corporate theory-in-use manifested itself in conflicts between divisional managers and central administration. These conflicts might have been the occasion for organizational inquiry and hence for double-loop learning. The parties to the conflict might have reflected on the conflict and its sources. They might then have explored and restructured the corporate norms, strategies, and assumptions which generated the conflict. None of this happened. Instead, the organization responded to the conflict, and to the inconsistency which underlay it, through a kind of single-loop learning, peculiar in the sense that the learners were not the organization as a whole, but groups and components within the organization. Thus, existing divisions learned to protect their territories more effectively and to become more mistrustful of central. NBD learned to seek business opportunities outside the scope of existing business. Top management learned to criticize particular failures of NBD, but to avoid criticizing NBD itself.

In this, the organization functioned as a collection of subgroups, more accurately, as an ecology of subgroups, since each group functioned within the environment made up of other groups and functioned itself as a part of the other groups' environment. Each group then set and solved the problems presented to it by its environment, inquiring at the level of single-loop learning. The net effect of this pattern of single-loop learning (which we call *ecological adjustment*) was to maintain the constancy of organizational theory of action. Requirements for development and for decentralized manage-

ment remained incompatible. Espoused theory and theory-in-use for corporate development remained incongruent.

It is not true, however, that organizational theory-in-use remained entirely constant. For each "success" in the realm of corporate development made divisional managers more mistrustful, more wary of central, more resistant to central's next incursion. It would be reasonable to predict that, in the future, such development successes would be few and far between.

The inconsistency in organizational theory-in-use had begun to take on the character of an organizational dilemma. That is, in the light of its conflicting requirements and its inability to resolve that conflict through inquiry, the organization had begun to find itself in situations of choice where all the options open to it appear equally bad.

What was it about the organization that created this dilemma?

One answer is that members of the organization had not reflected on and inquired into the issue. They had discovered neither the inconsistency in the requirements for development and decentralization nor the incongruity between espoused theory and theory-in-use for development. They had no map of the problem. In fact, they had treated the whole development process as undiscussable.

But this answer merely shifts the question. What prevented members of the organization from discussing the issue and mapping the problem? What prevented them, in short, from doing by themselves what the consultant enabled them to do?

Let us consider more carefully the consultant's process of intervention (we will be discussing it further in Chapter 9).

When the consultant found that different members of NBD held different and conflicting views of their problem, he brought them together to confront and discuss their differences.

Within the Mercury Corporation, the norms of the behavioral world induced members to express their diagnoses of sensitive issues in private, never in public. Public discussion of sensitive questions was considered inappropriate. It involved the risk of vulnerability to blame, and of interpersonal confrontation. Both were to be avoided.

As a consequence, the members of NBD did not know the extent to which they held different views of their own problem.

The consultant sought to test and explore the different diagnoses by collecting case histories of corporate development which were

perceived as successes and failures—that is, by constructing an organizational history of development.

Perceptions and memories of the development stories were scattered among individuals who held their views private. No one, for example, held in their mind the whole story of the resin/textile "success" because no one had experienced that episode from many sides and from beginning to end. The story had to be constructed by piecing together many scattered perceptions, and the norms of the behavioral world militated against such an enterprise.

The consultant urged the various members of the organization to work together at interpreting the meaning of the stories.

As long as the stories remained scattered and uninterpreted, the map of the development process remained vague and the diagnoses of the development problem remained ambiguous. But such vagueness and ambiguity were considered normal and appropriate, given the shared wish to avoid exhuming corporate failures (which might give rise to blame) and to avoid raising in public features of organizational life that were consensually treated as undiscussable.

The consultant, then, undertook and supported others in undertaking the public confrontation of different and conflicting views of the problem, the testing and exploration of differences through the construction of organizational history, the concerting of scattered perception of organizational events, the shared interpretation of data pertaining to development, and the joint mapping of the corporate development process.

All of these interventions had the effect of reducing certain *conditions of error* (scattered perceptions, vague maps, ambiguous diagnoses) so that still other conditions for error (inconsistency and incongruity in corporate theory of action) might be surfaced and subjected to inquiry.

In the language of the previous chapter, the consultant sought to facilitate a process of organizational double-loop learning by reducing the conditions for error which prevented shared perception of inconsistency and incongruity in organizational theory of action. But these conditions for error were reinforced by certain prevailing features of the organization's behavioral world—shared strategies in individual theories-in-use—which included the following:

- Let buried failures lie.

- Keep your views of sensitive issues private; enforce the taboo against their public discussion.
- Do not surface and test differences in views of organizational problems.
- Avoid seeing the whole picture; allow maps of the problem to remain scattered, vague, ambiguous.

But these strategies reflect deeper and more fundamental norms, strategies, and assumptions:

- Protect yourself unilaterally—by avoiding both direct interpersonal confrontation and public discussion of sensitive issues which might expose you to blame.
- Protect others unilaterally—by avoiding the testing of assumptions where that testing might evoke negative feelings, and by keeping others from exposure to blame.
- Control the situation and the task—by making up your own mind about the problem and acting on your view, by keeping your view private, and by avoiding the public inquiry which might refute your view.

These features of the behavioral world, which the consultant sought to by-pass and counteract, constrain and guide the character and extent of organizational learning. They determine what will be discussed and what will be left undiscussable, which individual perceptions of organizational experience will be left scattered and which will be concerted. They limit the extent to which organizational maps will be constructed, shared, and tested. They determine whether and in what way conditions of error will be reduced so that inconsistencies and incongruities in organizational theory-in-use may be discovered.

But these features of the behavioral world which constrict the possibility for joint inquiry into the problems of corporate development also enter into the corporate development process itself. Thus central administrators protected themselves unilaterally and sought to control divisional managers (by withholding information from them, for example), and divisional managers protected themselves and sought to resist control (by their territoriality and by their wariness of central). And both groups sought to avoid confrontation and the

evocation of negative feelings by avoiding the public testing of their assumptions about the other.

Thus features of the behavioral world which entered into the development process and helped to create its perceived difficulties also constrained inquiry into that process. One might say that the behavioral world protected itself from exposure.

The organization's theory of action is embedded in a behavioral world which shapes and constrains instrumental theory-in-use at the same time that it shapes and constrains organizational learning about theory-in-use. This is what we shall call the organization's *learning system*. Mercury's learning system limited organizational learning about corporate development to the process of ecological adjustment described above—that is, to a process in which each group carried out single-loop learning so as to cope, within unchanging norms, with the problems created for it by other groups in its environment. For a time, the consultant intervened so as to reduce conditions for error and to surface the central dilemma of development which the organization's learning had helped to create and preserve.

The Mercury story illustrates, then, some of the ways in which an organization's learning system may prevent double-loop learning, and it suggests a kind of limited intervention which by-passes features of the learning system so as to increase the likelihood of double-loop learning.

The Mercury story also suggests the dialectical nature of the larger process of organizational change within which we find episodes of learning or of failure to learn.

The organizational problems of corporate development arose and came to be perceived as a consequence of a dual process, each part of which took the form of the setting and solving an organizational problem. The phenomenal rate of corporate growth led to difficulties which were perceived as a problem of management overload, and were "solved" through decentralization. The wish to maintain that rate of growth, in the face of what was seen as declining R & D productivity, led to the creation of a New-Business Division. The working out of these two intertwined "solutions" led to the creation of a new inconsistency in organizational theory-in-use, an incompatibility between the requirements of decentralization and development which we have described as an organizational dilemma manifested in the conflict between central and divisional management.

We will use the term "*organizational dialectic*" to refer to such processes. In them, organizational situations give rise to organizational inquiry—to problem setting and problem solving—which, in turn, create new organizational situations within which new inconsistencies and incongruities in organizational theory of action come into play. These are characteristically manifested in organizational conflict. The organization's way of responding to that conflict yields still further transformations of the organizational situation.

We believe that organizational learning occurs within the frame of such dialectical processes, which stem from two conditions of organizational life: Organizations are necessarily involved in continual transaction with their internal and external environments (that is, in situations) which are continually changing both as a result of forces external to organizations, and as a result of organizational responses to their situations. Second, organizational objectives, purposes, and norms are always multiple and potentially conflicting.

As a consequence, it is no accident that organizational solutions give rise to further problems; they may be expected to do so, given the dialectic frame within which organizational inquiry occurs.

But this, then, raises sharply the problem of criteria for the evaluation of organizational change and learning. This problem is central to our inquiry, and will occupy us especially in Part III. At this point, however, we may note the following:

- The achievement of stable solutions is not an appropriate criterion for organizational learning; it is in the very nature of organizational problem solving to change situations in ways that create new problems.
- Organizational effectiveness—as measured by the achievement of espoused purposes and norms—is an incomplete criterion for organizational learning. It is appropriate in situations where error correction can occur through single-loop learning alone. It is insufficient where inconsistencies in organizational theory-in-use set requirements for double-loop learning.

"Good dialectic" is the term we will use to describe processes of organizational inquiry which take the form of single- and double-loop learning, as appropriate, and where (as we pointed out in the previous chapter) both single- and double-loop learning meet standards of high quality inquiry.

The achievement of good dialectic requires organizational deutero-learning. That is, it requires that the organization's members reflect on and inquire into their organizational learning system and its effect on organizational inquiry.

SUMMARY

The Mercury story illustrates some of the features of organizational learning systems which inhibit organizational double-loop and deutero-learning. The very features of Mercury's behavioral world which tended in the long run to make corporate development ineffective also worked to prevent members of the organization from collaboratively inquiring into the defects of the development process. The disposition to treat interpersonal and intergroup conflict as undiscussable, the *taboo* on public analysis of corporate failures, the wish to avoid direct interpersonal confrontation—all of these factors, and others related to them, contributed both to the ineffectiveness of the corporation's formal development agency and to the members' inability to diagnose and respond to patterns of ineffective development.

When, with the consultant's help, members of the corporation were able to piece together the stories of successful and unsuccessful development, their diagnosis revealed inconsistencies in organizational theory of action which had become organizational dilemmas. Imperatives of decentralized control led each baron to focus exclusively on the production of expected earnings and to cling to his own business territory; the imperatives of corporate development required the barons to view themselves as pieces of the corporate puzzle and to give up parts of their territories in the interest of the recombination and restructuring of existing business.

This dilemma expressed itself in a three-way conflict among the barons, central management, and the New-Business Division. Within the prevailing norms of the organization's behavioral world this conflict was undiscussable. Members of the organization could not, then, publicly reflect on the conflict in order to initiate a process of organizational double-loop learning. Instead, organizational learning took the form of ecological adjustment: Each unit of the organization learned (single-loop) to survive and to protect itself within the larger organizational context.

The organization had developed a pattern of intermittent corporate development which ran counter to its espoused theory of development and had little or nothing to do with the New-Business Division. But the organization had no map of this process; in a very real sense, it knew more than it could say. Moreover, the barons' increasing wariness of central management made this informal pattern of development less likely to succeed.

Within the prevailing learning system, members of the organization could not surface the underlying dilemmas which had arisen from the organization's earlier growth, nor could they surface the incongruity between their espoused theory and their theories-in-use for development. Mercury's learning system prevented its members from engaging in good organizational dialectic.

CHAPTER FIVE

A model of limited-learning systems

INTRODUCTION

We can now present a model of organizations whose learning systems are conducive to limited learning—a model of organizations which are unlikely either to correct first-order error by double-loop learning or to inquire into their own learning systems.

The model set out in Fig. 5.1 is a systems model which incorporates, in summary fashion, the information presented in the last two chapters. Beginning with primary inhibitory loops, the model traces their rings of consequences.

Primary inhibiting loops lead to secondary inhibiting loops—that is, to group and intergroup dynamics which reinforce conditions for error. Within such processes, some kinds of error remain correctable, while others do not.

Correctable errors enter into a learning cycle which eventuates in action which, in turn, yields new error or appropriate response (a match of outcome to expectation).

Uncorrectable error tends to be camouflaged (that is, to be hidden, denied, or disguised). The camouflage of uncorrectable error leads to new primary loops which make organizational double-loop and deutero-learning unlikely and, for individuals, make double binds likely.

Figure 5.2 (pages 112, 113) displays a more complex version of this model. Reverse arrows along the bottom of the figure indicate feedback loops that close the system. The order of columns, from left

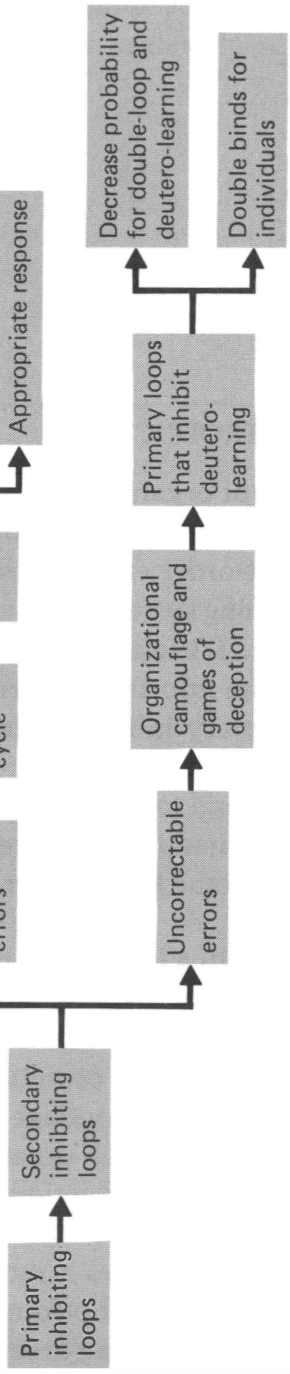

FIG. 5.1
Model O-I: Limited Learning Systems

to right, and the numbered arrows in the reverse direction, show the interaction effects which seem to us to be most important. Complex as it is, the model is still oversimplified in a number of ways. For example, arrows along the top might also have been numbered; secondary loops, to take one instance, lead not only to correctable and uncorrectable errors (Col. 7) but also to camouflage of error (Col. 8). Each column has its effect not only on the one immediately following but on others further down the line. However, we have tried to arrange the columns so that left-to-right order is a pretty good representation of direct effects.

The model does not describe the etiology of limited-learning systems. Its meaning is *not*, for example, that primary inhibitory loops came first in the evolution of organizations and that they led later on to dysfunctional group and intergroup dynamics, which still later led to deutero-learning loops. We think it more likely that a limited-learning organization, at any period of its evolution, displays at least embryonically the full configuration of the system.

What the model does reveal is the set of direct and indirect effects and feedback loops which interconnect with the principal elements of a limited learning system. Given any column, such as Col. 8 (camouflage), one can look to the left to find its immediate and less immediate conditions, and to the right to find its immediate and less immediate consequences.

We have begun with primary inhibitory loops because they seem to us the best starting point in order to explain a limited-learning system. In addition, as we will argue in Part III, they seem to us the best starting point for intervention. They are "primary" not in the sense of temporal order, but in the sense of their primary importance among the processes which make up the system.

The model, then, has the principal function of guiding the mapping and diagnosis of limited-learning systems. Because such mapping is essential to effective intervention, as we understand it, the model will reappear in that context, as well.

We will proceed through the model, from left to right, taking up each column and its interactions with others.

COLUMN 4: PRIMARY INHIBITORY LOOPS

These are the loops described in Chapter 3, and illustrated in the cases of Carlos, Roberto, the Principal, and the Finance/Line officers.

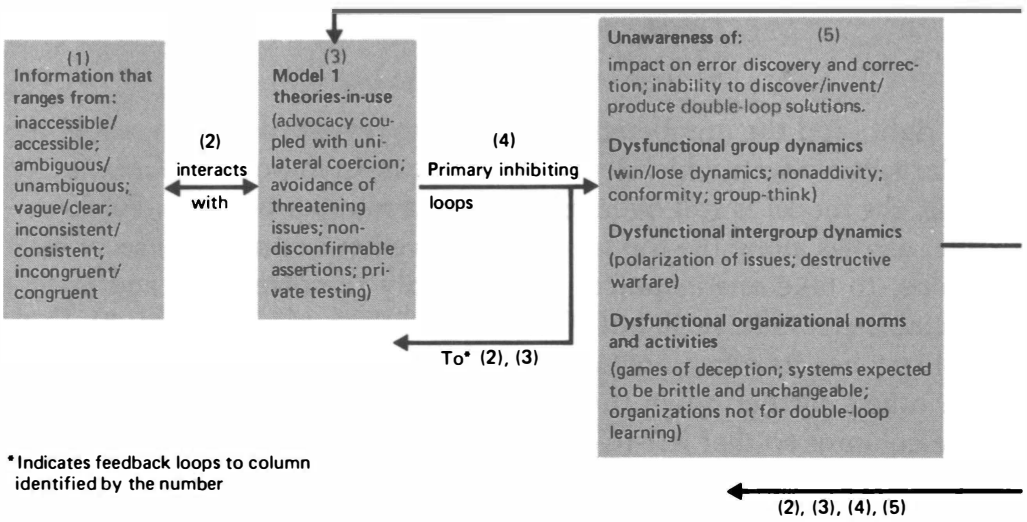

FIG. 5.2
Model O-I: Limited Learning Systems

Elements of an organization's instrumental theory of action are inaccessible, unclear, or inadequate. One or more of these features of organizational theory of action gives rise to error (Col. 1). In a good dialectic, such conditions of error would be confronted and reduced through organizational inquiry. In a Model I behavioral world, however, such conditions trigger Model I interactions (Col. 3) which reinforce those conditions for error, or create new ones.*

In the case of the government agency, for example, ambiguity over decisions which may or may not have been made induces some staff members to decide privately that "I am always involved in a discussion, never a decision," which further contributes to the ambiguity over decisions.

Within such loops, conditions for error become uncorrectable and trigger the very responses which make them so.

COLUMN 5: UNAWARENESS; DYSFUNCTIONAL DYNAMICS

Primary inhibitory loops reinforce unawareness of their effects on organizational learning.

Carlos, for example, is aware that he has failed to find a good solution to the problem of project assignment. He is unaware, how-

---

* "Model I" refers to *personal* theories of action and to a behavioral world made up of interpersonal interactions. "Model O-I" refers to an organizational learning system.

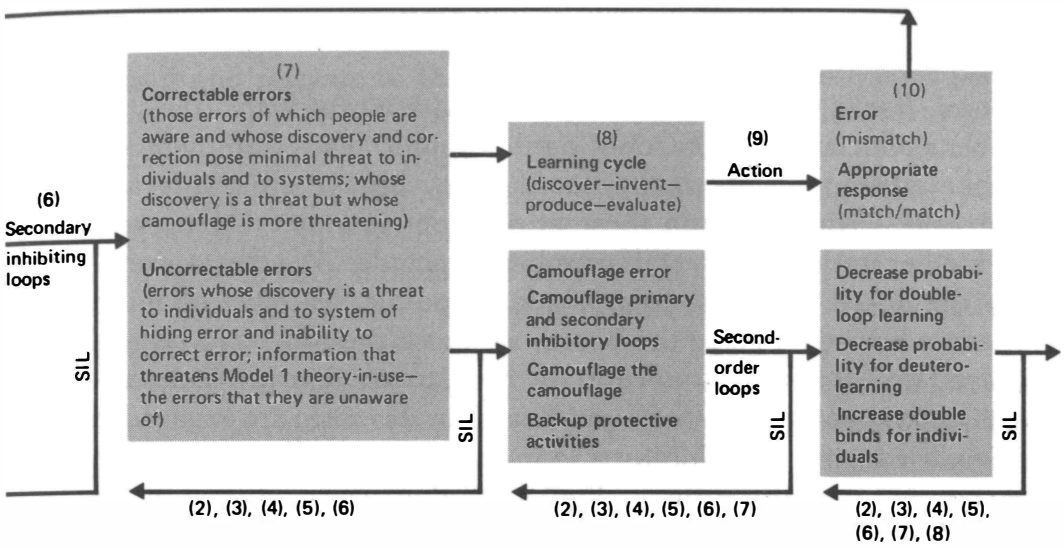

ever, how he has contributed to interactions which perpetuate the vagueness of criteria for project assignment, thereby increasing the likelihood of continued failure. Nor is he aware of his inability to design and contribute to interactions which would surface and confront both conditions for error (such as vagueness) and the Model I behaviors (including his own) which reinforce those conditions.

Given such awareness, Carlos would be more likely to seek to reduce his own contribution to the primary loops. Without such awareness, he is unlikely to do so. Hence his unawareness helps to preserve the primary loops. But the primary loops tend to reinforce his unawareness. Within the loop described in his case, for example, Carlos blames his failure on the Program Manager or on the unreasonable behavior of staff members. He sees the ambiguity of role definition, in which he and the Program Manager are both embroiled, as the latter's failure to take a firm stand. Such perceptions, sustained by the primary loop, are incompatible with awareness of the loop itself, Carlos's contribution to the loop, or the conditions for error reinforced by the loop.

Primary inhibitory loops yield intra- and intergroup dynamics which mirror and amplify the properties of primary loops. These secondary loops feed back, in turn, to sustain primary loops.

In the case of the government agency, a particular interaction between a staff member and members of the Board (Board refusal to allow a staff member to participate in discussion of the staff member's

proposal, Board rejection of the proposal without feedback, the staff member's reluctance to confront Board behavior) leads staff members to make attributions about the Board ("they behave arbitrarily and unilaterally") and to act on those attributions in ways that prevent their further testing. Thus, staff members conclude: "We must second-guess the Board, we must submit safe proposals, we must establish consensus before we go to the Board, we must 'prep' so as to be well buttoned up." It becomes appropriate, then, to say that whole departments of an organization (Staff and Board, in this case) are engaged in an amplified primary inhibitory loop in which, for example, each makes untestable attributions about the other, each induces in the other the behavior most feared and disliked, and both contribute to the conditions of error (vague project criteria, ambiguity over decisions) which trigger such interactions.

Within such an intergroup context, small-scale interactions are still more likely to take on the properties of primary loops.

Sustained primary loops lead to the expectation that organizations are brittle and unchangeable. Where members learn to despair of double-loop learning, the stage is set for games of deception.

These are conditions for intergroup interactions of the sort described in the case of higher education. The presidents of public and private universities, having presented themselves to the public as engaged in an effort at cooperation, find themselves unable to cooperate. Their response is to engage in debate with one another, in private complaints, and (in this case) in leaks to the press, where each party accuses the other of betrayal and seeks to impose on the other the burden for failure of cooperation. The process takes on the character of a game in which winning consists of minimizing one's own risk, forcing the other to take risks, avoiding one's own responsibility for failure, and putting blame for failure on the other.

In Roberto's case, the conditions were created for another such game of deception. Roberto informs us that he and the Director eventually worked out a solution in which a new man would be hired, nominally reporting to the professor, but in fact taking sole responsibility for the use of the equipment. A necessary condition of the arrangement is that all parties agree to deny the arrangement. Now, if a client were to complain about performance on a project, the conditions would be set for the new man and the professor each to avoid for

himself and relegate to the other the responsibility for poor performance, all the while refraining from public acknowledgment of the actual working relationships.

Games of deception, of gaining credit and avoiding blame, have a tendency to occupy the foreground of organizational attention. They loom large in each individual's universe of concerns, distracting him or her from awareness of the uncorrectable errors and the related processes which underlie it.

Moreover, such games create an impression of organizational fragility and rigidity. Polarization of groups and persons allows each person and group to feel that it is blocked by others. Each member of the organization, aware of the layers of potential vulnerability shared with others and of the games designed to protect against that vulnerability, experiences the organization as brittle. A false move, an unwitting disclosure, a direct confrontation, and the house of cards might come tumbling down. It is inherent in such judgments that they are unlikely to be put to the test.

COLUMN 7: CORRECTABLE AND UNCORRECTABLE ERRORS

None of the processes so far described prevents members of an organization from detecting and correcting errors in first-order performance, so long as the detection and correction do not require that the conditions for error be confronted in ways that threaten Model I governing variables.

Thus members of the Mercury Corporation, as described in Part I, could detect and respond to routine errors in production of quality control. It is also entirely plausible that the Principal and the third-grade teacher could jointly detect and correct errors in the assignment of students to classrooms.

But Carlos and the PM could not correct errors in task-assignment. Members of the Agency could not correct for delays in proposal development. The Chairman and Chancellor of the Board of Higher Education could not remedy the failures of public/private cooperation.

Given the primary and secondary loops characteristic of these living systems, such errors become uncorrectable. Given the frame of conditions for error and Model I theories-in-use, efforts at error-

correction tend, in fact, to amplify errors. In the government Agency, the act of "putting one man in charge of a project" would simply "put us right back in the glue." The Chairman and Chancellor's resort to private meetings simply exacerbated the Forum's climate of distrust.

The constraints imposed on organizational learning by limited-learning systems depend on the scope of uncorrectable error. We have argued that errors tend to be uncorrectable when their correction would entail double-loop learning for the organization; that is, when norms central to organizational theory-in-use would have to be questioned and changed. We have also argued that errors tend to be uncorrectable when their correction would threaten Model I governing variables; that is, when it would require double-loop learning at the level of the behavioral world. Within the range defined by these criteria, however, there is room for great variation. And organizations vary greatly along these lines, as we will see in the chapters that follow.

COLUMN 8: CAMOUFLAGE

In a Model I behavioral world, the discovery of uncorrectable errors is a source of personal and organizational vulnerability. The response to vulnerability is unilateral self-protection, which can take several forms. Uncorrectable errors, and the processes that lead to them, can be hidden, disguised, or denied (all of which we call "camouflage"); and individuals and groups can protect themselves further by sealing themselves off from blame, should camouflage fail.

Camouflage may take the form of resort to espoused theory ("We are open, trusting, and cooperative with one another"), where everyone makes an open secret of the incongruity. Or the uncorrectable error may be attributed to external factors, over which members of the organization have no control. Members of the organization may make a public show of attack on the problem while sharing an understanding of the ritual nature of that attack.

Protection often takes the form of anticipation of the consequences of uncorrectable error so as to give the anticipator a margin for acceptable performance. Staff members in the agency, for example, might deliberately overestimate the time required to bring a proposal to fruition in order to compensate for the delays they have learned to expect. Or one might cover oneself by documenting his or

her own performance in such a way as to throw the blame for delay upon others. These are phenomena familiar to all members and students of organizations.

What is often less familiar are the consequences of camouflage and protection. Camouflage is a response to uncorrectable error which draws off energy which might be used to engage such error. The hiding, disguising, and denying of uncorrectable error tend further to protect it from inquiry and thereby to reinforce second-order loops.

Moreover, when camouflage and protection are broadly practiced, they set the conditions for a second layer of camouflage. The hiding, denial, or disguising of uncorrectable error cannot come to light without actualizing this double layer of vulnerability. Hence, *these* processes must be hidden, denied, or disguised.

In the case of the agency, there is as much of a taboo against public discussion of the fact that we are making the "narrow funnel" an open secret as there is on discussion of the "narrow funnel" itself. In effect, we must tacitly agree not to discuss our denials and disguises if they are to do their job. This, then, represents a further impediment to deutero-learning. The very allusion to such defenses may create a sense of intolerable risk.

COLUMN 9: PRIMARY LOOPS THAT INHIBIT DEUTERO-LEARNING

These are primary loops which arise in the course of attempts to inquire into an organization's first-order activities. They are generated by the same kinds of factors—conditions for error, Model I behavioral world—which create first-order primary loops. But they are also reinforced by the secondary consequences of first-order loops. And they feed back to reinforce both of these.

In the case of the Agency, the very existence of a problem in first-order performance is ambiguous. ("Do the delays in launching the new Institute constitute a problem, or do they not?") This ambiguity is sustained, in part, by the Associate Director's denial of the problem and by staff members' reluctance to confront that denial. But it is also sustained by the image each person holds of the prevailing norms for Staff/Board and Staff/Director interactions, norms which stress avoidance of confrontation, open secrets, and incongruity. In their collusion to maintain the ambiguity of the problem, members of the

organization need have no sense of extraordinary experience; they are simply conforming to the norms for intra- and intergroup behavior.

When such normal behavior prevents inquiry into the learning system of the organization, however, it serves to maintain the prevailing patterns of the learning system. The primary loops that inhibit reflection on the learning system protect the complex of primary and secondary loops that make up the learning system.

COLUMN 10: DECREASE PROBABILITY FOR DOUBLE-LOOP AND DEUTERO-LEARNING; INCREASE DOUBLE BINDS FOR INDIVIDUALS

Double-loop learning depends on awareness of error, which primary- and deutero-learning loops prevent. When errors are uncorrectable, they cannot trigger double-loop learning.

Members of a limited-learning system might inquire into the features of their system which make errors uncorrectable. But deutero-learning loops prevent such inquiry. Hence, in limited-learning systems, double-loop and deutero-learning are unlikely.

From the viewpoint of an individual living in such a system, the organizational world is apt to be peculiarly frustrating and constraining. Such a person is apt to find himself or herself in lose/lose situations which present intractible dilemmas. A staff member of the government agency, for example, is likely to experience a world of ambiguity in which he fears punishment for delays he does not see as his responsibility; fears the consensus-seeking process and fears to avoid it; feels he does not understand the problem; and believes he cannot raise with others the issue of the factors that have put him in this situation. The school principal finds herself ineffective in her efforts to get the teacher to recognize his poor performance, and feels vulnerable as a consequence; but the prospect of a more direct confrontation with the teacher makes her feel equally vulnerable. The resulting dilemma is one that she also feels unwilling to surface, since it could be taken as a sign of poor performance which would bring with it still further vulnerability.

These are situations that meet the conditions Gregory Bateson has laid down for "double binds" (Bateson, Gregory, 1972)—namely, one is caught in a no-win game and the rules of the game are undiscussable.

Limited learning systems are predictable generators of situations such as these. They require, as a condition of membership, that individuals assume the double layers of vulnerability inherent in camouflage and games of deception. They then put a taboo on discussion of these conditions and on the processes by which one has gotten caught in them.

Bateson advances his theory of the double bind to account for the etiology of schizophrenia. Double binds often reinforce features of limited learning systems by intensifying Model I behavior, by generating internal stresses conducive to deteriorating performance or to sporadic eruption, and by creating double-binds for others.

SOURCES OF THE O-I MODEL OF ORGANIZATIONS

The O-I model of organizations—the model of limited-learning systems—is so awe inspiring in its apparent irrationality that one is moved to ask how it can possibly result from apparently rational processes of organizational design. After all, formal organizations are consciously designed; they are calculated to achieve intended objectives. But when we look at what actually goes on in organizations, as we have been doing throughout Part II, we find much that is counterproductive to the original design. Moreover, the counterproductive activities are just as obvious and commonplace as the ones that are more nearly congruent with the original designs. How is it that people design organizations in ways that eventually allow or induce these organizations to take on the characteristics of limited-learning systems?

In order to answer these questions, we should remind ourselves that to design an organization is to calculate ahead of time what the organization needs to accomplish its tasks. In order to make these calculations, the designers must have some model of where they are, where they are going, and how they are to get from where they are to their destination. Obvious as this statement may be, we know very little about these calculated design or map-making processes. As we shall see, one of the great gaps in our knowledge is in understanding the causality of the obvious.

When people begin to design an organization, they are usually aware of some objectives for the organization and the resources that will be available to it. They also know something about what they can

or cannot do with these resources. For example, they know something about how money may and may not be used, policies that are legal or illegal, time perspectives that are realistic and unrealistic. The constraints embedded in these resources are relatively clearly defined and rigid. This is true for human beings as resources. We know very little about their potential limits, and what we do know suggests that, for a fair reimbursement, human beings are willing to be quite flexible and work under less than optimal conditions.

There is one quality about human beings, however, that is crucial and not as flexible as many other human qualities. We are speaking of the way human beings process information and think; the way they deal with complex problems. It appears that people think hierarchically when they deal with complexity (Miller 1956; Simon 1968 [as already noted]). For example, a person may have some experience with and remember different automobiles. As the number becomes large, there is a need to store the information effectively so that it can be retrieved efficiently. One device is to create higher order constructs which permit the individual to subsume many of the individual cars. The constructs may be in terms of color, size, or any other attribute. But whatever attribute has been selected, it becomes the key to dealing with the cars. If the numbers continue to expand, the individual may invent new and even more abstract categories. Again the invention means that some new attribute has been selected in order to remember and to form the basis of action. The basic assumption is that whatever the attribute, it is adequate for storing and retrieving that information.

We suggest that these basic thought processes are also used in the design of organizations. Organization designers begin with people who are to act as agents for the organization. To be an agent of an organization means to produce that behavior which the organization is thought to require. And this behavior, under prevailing organizational paradigms, is usually conceived as a patterning of components that looks like Fig. 5.3. The pyramidal pattern, which is so invariant that people often define organizations as pyramidal structures, derives from the principles of specialization of work and of hierarchical control.

Designers think hierarchically about organizations, just as they think hierarchically about complex problems in general. Complex tasks are to be broken down into relatively simple ones and grouped

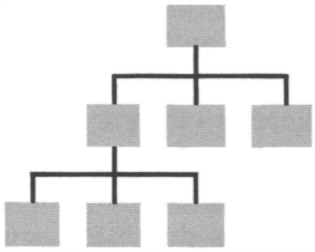

FIGURE 5.3

for ease of recognition and control. This strategy, which suits the pattern of hierarchical thinking, also suits the designers' wish to mini-max, that is, to devise task systems which minimize the organization's loss and maximize its gains. For designers tend to mini-max by defining roles that almost anyone can fulfill. The resultant specialization of work, with its associated pyramid, appears to the designers to make it easier to find people, educate them to the requirements of the organization, and manage them.

But in using this strategy the designers have created two major new problems. The first is an organizational problem. If people specialize, then their efforts must be coordinated and controlled in order to achieve the desired goals. In order to coordinate and control, there is a need for people to perform these functions. Since these people, who are assigned the roles of superiors, are human beings, there are limits to how much information they can process; and so there are limits to the number of people who may report to them.

Also, the designers are aware that people are organisms with their own motivations, theories of action, and goals. Therefore, they try to give the superiors the resources they need to minimize complexity and noise. For example, the superiors are given power to control and coordinate; to reward and penalize; to evaluate performances, etc.

The combination of specialization of work and centralization of power and information (which when diagrammed is the pyramidal structure) has the advantage of making organization activity manageable. However, built into this design are some contradictions that create tensions and discontinuities that must be dealt with.

Let us begin with the inner contradictions that arise from the fact that this mini-max strategy is being used on individuals. They may be summarized as follows.

1. People in our society are programmed with Model I theories-in-use. The bases for competence and a sense of effectiveness are behavioral strategies of advocacy coupled with unilateral control over others, over the immediate environment, and over the task to be accomplished. But these are the very conditions that are taken away from them by the organizational designers in order to mini-max. More accurately, the lower one goes down the organization, the less it is possible for the people to have these conditions.

2. People in our society are educated to develop many and complex abilities. Yet, the lower they go down the organization, the less it is likely that they can use these abilities.

Until recently, the most frequent strategy for reducing these inner contradictions was for organizations to pay as fair wages as they could, to guarantee (with the influence of the unions and the government) as much job security as they could, and to strive to educate superiors so that they did not use leadership styles (e.g., authoritarian, unilateral) that compounded the problem.

These strategies have probably worked as effectively as was possible, given the constraints imposed by human beings. For example, a worker may be willing to be paid off to forget that he is composed of more than the hand the company wishes to hire, but to do so he will have to learn to deal with boredom. One can suppress feelings of boredom only to a certain limit, so employees create all sorts of informal activities that liven up their world, especially such activities as breaking rules.

Workers may also be willing to give up much of their control over their work world if they can trust their management to be fair and not take advantage of them in their dependent, submissive position. But the trust between management and the worker is difficult to develop and even more difficult to maintain. The same inquiry suggests that both employees and management may be responsible for these difficulties.

3. People in our society are programmed with Model I theories-in-use that predispose them toward single-loop learning rather than double-loop learning. The result of this is a paradox. On the one hand, organizations prefer people who remain at the single-loop level of learning, especially if they accept the game plan of the organization. This makes it possible for organizations to accomplish some of their

most important functions; i.e. to create continuity, consistency, and stability, and to maintain the status quo in order to achieve objectives within desired costs.

On the other hand, organizational designs are imperfect and incomplete. They require continual reflection and monitoring to meet challenges from the changing external environment as well as from the counterproductive activities of the internal environment. Organizations must be concerned with discontinuity, inconsistency, instability, and changes in the status quo. Otherwise, they take on the quality of dynamic conservatism, a quality that may threaten long-range survival. These conditions require double-loop learning, a competency that we have just said people do not tend to develop in our society.

The problem becomes even more complex because people with Model I theories-in-use will not only tend to create ineffective problem solving, they will tend to create defensive group dynamics. For example, group processes tend to become highly competitive, rarely additive, low in openness, trust, and risk-taking, high in closedness, mistrust, and emphasis on not rocking the boat. Group-think, identified by Janis (1972), is a natural consequence. Organizations are therefore full of intergroup warfare and rivalries.

To compound the problem even further, since most of these consequences are inimical to the organization's effectiveness, those that carry them out expose themselves to potential punishment. Organizations are full of rules and policies that caution people against violating regulations and identify penalties for costly errors.

The difficulty, however, is not that policies exist and penalties are known. The difficulty is that to overcome these consequences requires reflection and learning of the double-loop variety—and that, we suggest, is inhibited by individual lack of competence reinforced by defensive group and intergroup dynamics. The problem is not that these conflicts and rivalries exist; it is that they are not discussable. It is not possible to manage effectively events that are not discussable.

Moreover, the fact that everyone knows that these counterproductive activities exist and are not discussable means that people are, in effect, in collusion to deceive the organization and each other.

How do people cope with deception that appears to them to be inevitable and necessary? One way is to develop games which are

played in order to maintain open secrets, avoid punishment for error, and camouflage uncorrectable error.

One consequence of gamesmanship and mistrust is that executives place greater emphasis on the use of rationality, direction, control, rewards, and penalties. In practice, this tends to mean that they begin to check on other people's work, not only to see if it is done, but also to see how it was accomplished. They also operate by asking detailed questions about issues and problems that may exist at levels lower than that of the person being questioned, but for which he or she is responsible. For example, they may ask a personnel vice president for the capacity of a parking lot in a plant away from the home office.

The result of such action by the superior is to create defensiveness in the subordinate. The subordinate now checks constantly on all details so that he or she will not be caught by the superior. However, the activity of the organization is not carried forward by such behavior. The result is simply to make the subordinate (and usually his or her subordinate) more defensive. Their response is to build up organizational defenses to protect themselves.

Organizational defenses may therefore be developed to protect various individuals and groups. These defenses can be used to needle people, which tends to occur when the rational methods seem to fail. But since the use of feelings is deviant behavior, and since neither the superiors nor their subordinates have much experience in their use, the tendency may be to create feelings that are much stronger than the situation warrants.

Executives may speak of "needling" the boys, or "raising a little hell to keep them on their toes," and so on. But if these conditions continue, it is not long before the "hot" decisions of the organization are administered by using emotions. This is commonly known in industry as management by crisis.

As management by crisis increases, the subordinates' defensive reactions to the crises will tend to increase. One way for subordinates to protect themselves is to make certain that their areas of responsibility are administered competently and that no other peer executive "throws a dead cat into their yard." Thus each subordinate's attention will become centered on the interests of his or her own department—and as department-centeredness increases, interdepartmental rivalries will tend to increase. The result is to decrease the organiza-

tion's flexibility for change and the cooperation among departments. In turn, top management will tend to adapt to this decrease by increasing directives, which again begin to recentralize the organization.

The external commitment, conformity, interpersonal mistrust, ineffective decision making, management by crisis, and organizational rigidity will tend to feed back to reinforce each other and to decrease interpersonal competence. Moreover, each will feed on all the others to reinforce itself. We would conclude that, under these conditions, the tendency will be to increase the energy required to produce the same input, or someday it may decrease the output, even though the input remains constant. When this state of affairs occurs, it may be said the organization's effectiveness has begun to deteriorate.

As the level of problem-solving effectiveness goes down, confidence in the processes also goes down. Feedback from a decrease in confidence will make future problem solving less effective, and this will require further counterproductive activities, which in turn will require more games. Thus we have several levels of circular processes serving not simply to maintain the counterproductiveness of the status quo, but actually to increase it.

So much for the inner contradictions caused by using human beings as agents. We now return to the inner contradictions created by the very nature of organization. As we have suggested, the basic purpose of organization is to maintain stability in order to achieve its goals and to survive. But in social organizations the requirement of double-loop learning and change becomes especially important, because the external environment tends to be dynamic and changing and because the internal environment is also basically unstable.

Thus social organizations have some built-in contradictions that arise from the paradoxical requirement of maintaining yet changing the steady state. The result is that agents are required to think and behave in contradictory manners. For example, agents may be told:

1. Take initiative.	1. Don't violate rules.
2. Sound alarms early for errors.	2. You will be penalized if errors are made.
3. Think beyond the present.	3. It is the present performance that is the basis for rewards and penalties.

4. Think of the organization as a whole.	4. Don't cross into others' areas of responsibility.
5. Cooperate with others.	5. Compete with others.

Again, the troublesome fact is not that these inconsistencies exist, but that they are not discussable. The personnel involved are not competent to discuss the undiscussable, nor do group processes and intergroup norms encourage them to develop that competence. Again, these conditions will lead to a set of games to protect people who must behave in counterproductive ways if they are to survive—or indeed, if the organization is to survive.

Earlier we noted that many of the activities "tacked on" to the original organizational design were counterproductive. Now we realize that they may be counterproductive but rational, given the fact that designers seek to minimize costs and maximize gains by using the behavioral strategies embedded in the pyramidal structure.

Participants involved in organizations, and scientists trying to comprehend them, cannot ignore these "added" activities. Both the original design and the added parts, within the context of the behavioral world, make up the larger, more comprehensive whole that we have called the learning system. For example, quality-control departments and budgeting departments may be added to control for the counterproductive tendencies of employees at the lower levels. Management information-systems departments may be created to control for possible counterproductive activities at all levels of the organizations.

The most important and least consciously and calculatedly recognized parts of the learning system are the "informal" systemic activities, such as rate-setting through group pressure, distortion of reality so as not to upset top management, intergroup administrative warfare, norms of psychological withdrawal, competitive win/lose dynamics, etc. Assuming that the original system was competently designed, if that system does not work effectively it would be because (1) parts of the learning system have evolved *and* (2) most of these parts are publicly known yet covertly maintained; they are omnipresent in everyday problem solving yet consciously not discussed with the upper levels of power.

The reason that we give great emphasis to these parts is that they are typically underemphasized. Indeed, they are suppressed, yet they

are critical for helping an organization to move toward double-loop learning. No matter how complete it may be, a diagnosis that is limited to the formal organization will not help the people to understand what the undiscussable subjects are, why they are undiscussable, or how to learn to discuss them.

One cannot understand the present state of an organization merely by knowing its original state. The parts of an organization are multilevel and more complex than the original design. To understand their existence and their interrelationship one must first observe the inner contradictions of formal organizational design. Then one must recognize that, while human beings may be willing to fulfill the organization's initial demands, they must also come to grips with the problems that they and their groups and intergroups create for the organization and for one another.

People who live in the learning systems of most organizations come to conceive of the factors that inhibit double-loop learning as natural and expectable. They do this for sound reasons. In a world where double-loop learning is not sanctioned, it *is* natural to expect double-loop learning to be unlikely. Similarly, people come to believe that organizations are brittle and hopelessly unchangeable. People become upset, for example, if they are asked to focus on deception and mistrust. They do not know how to deal with such feelings. Hence, the feeling of brittleness.

Undiscussability, rigidity, and emphasis on the status quo come to be seen as part of common sense, and, as Geertz (1975) has shown, thereby acquire additional qualities of naturalness, practicality and embeddedness. Through processes such as these, the "rational" patterns of organizational designers come to resemble limited-learning systems.

PART III

INTERVENTION TOWARD O-II ORGANIZATIONAL LEARNING

In Part III we present a model of an organization learning system (O-II) that can be used to decrease the inhibitions to double-loop learning embedded in Model O-I. We also include several probable scenarios of how organizations would go about double-loop learning.

Next, we address the question, how do we get from here to there? How do we help organizations move toward Model O-II learning systems? The answer to this question is, in effect, our model of intervention. We differentiate between limited and comprehensive intervention and note the conditions under which each may be utilized.

CHAPTER SIX

O-II learning systems

An O-I learning system cannot learn to alter its governing variables, norms, and assumptions because it does not permit double-loop learning to take place. If we are interested in overcoming the forces that inhibit double-loop learning, we must seek another learning system. But if our assertion that most organizations contain O-I learning systems is valid, and if Model O-I systems deter the creation of new learning systems that run counter to their basic structure, then the new learning system we seek is not likely to be found by looking at the world as it presently exists. The creation of Model O-II will therefore have to be a rare event.

Rare events cannot be created without a map that describes the new territory. A map is needed of a new learning system that provides us with a picture of the probable end state. Our candidate for this new learning system is Model O-II. But if O-II learning systems can not evolve from O-I, then we need a map of how to move Model O-I toward O-II. Such a map should inform us of the conditions that, if violated, would take us off our course. For example, if we make non-disconfirmable statements, or if we unilaterally control others in order to win, then such behavior will inhibit the likelihood of achieving Model O-II. Models II and O-II can be used as guides to the conditions that should not be violated during the transition.

However, these models are not sufficient for the guidance needed. Rules, maxims, heuristics are needed to guide people in advocating a

position and coupling it with inquiry. Also rules are needed to inform the individuals how to make disconfirmable statements. These rules are the operational definitions for action without which the guideposts provided by Models II and O-II cannot be achieved.

One other caution. When we speak of Models II and O-II as being end states there is the risk of implying that there are final states which, when reached, are fixed and unchangeable. There are at least two reasons why they are not. The first is that Models II and O-II represent ideal states that may never be achieved but only approximated. Such maps are valuable to have because they provide models for creating a good organizational dialectic. This leads to the second and more important reason. Models II and O-II will not tend to become fixed and rigid because of their built-in capacity for double-loop learning. Such learning continually questions the status quo.

Since O-II learning systems are rare phenomena, it will not be possible to provide as rich descriptions of actual examples as was possible for O-I learning systems. Neither of us knows, for example, of an organization that has a fully developed Model O-II learning system, nor are we aware of any such O-II learning system described in the literature. Indeed, we believe that the intervention theory needed to transform organizations to include O-II learning systems is also extremely primitive. The best that we are able to do is present cases of the beginnings of Model O-II learning systems in various settings in which we have worked.

CARLOS AND THE PM: A NEW SCENARIO

A key component of an O-II learning system is the theory-in-use which people use to deal with the conditions of error (i.e., scatteredness, ambiguity, vagueness, etc.). O-II learning systems require conditions (Chapter 3) under which mistaken assumptions can be reformulated, incongruities reconciled, incompatibilities resolved, vagueness specified, untestable notions made testable, scattered information brought together into meaningful patterns, and previously withheld information surfaced. As we have suggested, these conditions are highly unlikely when people utilize Model I theories-in-use because they create primary inhibiting loops. What would an individual theory-in-use look like that helped people to confront conditions of

error in ways that led to the conditions being reduced and the error corrected?

Let us return to the case of Carlos. Let us suppose, contrary to fact, that he and the PM set about trying to confront and remedy the first-order errors in organizational theory-in-use that gave rise to their uncomfortable interaction. What would their inquiry have been like? What tasks would they have had to perform, and what difficulties would they have encountered?

Carlos might have begun by saying: "We have a request from Company X to carry out an audit of new product possibilities, and I am unsure how best to assign the task. But you know, I sense an additional difficulty. It is not clear to me how you and I should be working out problems such as these."

Carlos might then have said that he had wanted the PM to suggest a way of proceeding in the Y case—but he wanted also to be able to question suggestions of that kind so as to be able to understand the reasoning behind them. Perhaps the PM had interpreted his questioning as a rejection of advice, but this was not what Carlos had intended. And Carlos, feeling that the PM would not come forward with a proposal, had assumed that *he* must take initiative. This he had done. He had been unhappy with the results. But he had assumed, on the basis of his earlier attribution to the PM, that the PM did not want to take responsibility for engaging matters of this kind, and so Carlos had kept his unhappiness to himself. As a result of the Y incident, he had wondered what the original intentions had been in setting up counterpart managers. If there were policies for the interaction, Carlos did not know what they were. But Carlos had not felt free then to raise this question with the PM; after all, the PM had given no sign that he thought the relationship needed discussion.

In such a discussion, Carlos and the PM would be attempting to discover the nature of their situation. They would be examining their respective pictures of the policies that governed their relationship and of the events that had actually transpired. As each revealed his interpretations of those events, the assumptions he had made, and the ways in which interpretations and assumptions determined his subsequent behavior, he would be providing data against which the other could test his interpretations and assumptions. In such an inquiry, each would depend on data provided by the other. Together, they

might then be able to construct a picture of the pattern of interaction into which they had fallen. They would be mapping the small-scale learning system they had been in process of creating with and for each other.

PM: What do you mean?

Carlos: How should we be dividing up the responsibility for making this kind of decision? In past cases, like the Company Y request, I felt you were giving the ball to me. I accepted it then, even though I felt unsure how to proceed and would have welcomed advice from you.

If the PM had felt nondefensive and free to enter into this inquiry, he and Carlos might then have reflected on the case of Company Y. How did each of them perceive what went on in that case? What did the PM assume about Carlos, and Carlos about the PM, and how did each behave on his assumptions? And how was that behavior then perceived by the other? The PM, for example, might say that he saw himself primarily in a teaching role here; Carlos was, after all, his counterpart and would eventually take on the program management job. He (the PM) was unsure how best to do this teaching. And there were, so far as he knew, no policies for it. He had at first tried to offer advice, but had found that Carlos, in the Y case, rejected advice and went on to do it his own way. Though offended by this at first, the PM had decided that it might be best, after all, to let Carlos try his own approaches and make his own mistakes. Perhaps he could teach best by giving Carlos a free hand and by remaining available for consultation if Carlos wished it.

Carlos might have been surprised at the PM's perception that his advice had been rejected. He might ask the PM for his memory of the events that led him to this interpretation. The two men would then compare their memories of events, and test their differing interpretations against those events. They might or might not be able to resolve that difference. They might, in any case, notice that their differing, private interpretations had gone unexpressed and untested.

They might then realize that, together, they had generated a behavioral world, self-sealing up to that point, in which each had been responding privately to the uneasiness generated by the ambiguity in their relationship. Carlos had expected and wanted initiative and advice from the PM and had been blind to the ways in which he

(Carlos) showed himself unwilling to accept such advice. The PM had privately interpreted Carlos's behavior as a sign that Carlos wanted to be left alone, and, again privately, had revised his interpretation of his own teaching role according to this assumption. Each privately attributed to the other the responsibility for initiating the present pattern of their interaction.

Such a picture is at once a map and a diagnosis. It draws on remembered history to construct a model of existing patterns of interaction which, in turn, account for what each finds troublesome in the situation.

But the use of such a map/diagnosis is that it gives Carlos and the PM a way of thinking about how they might now redesign their pattern of interaction, on the occasion of the present incident, the Company X request. This requires invention. What is to be invented is a new pattern of interaction for dividing up and sharing responsibility for decisions like the present one. Not only "new" but "better"; "better" in terms of their respective expectations for the relationship, and their respective images of the kind of interaction that will be organizationally effective. But then, the meaning of "better" will have to become a part of their shared inquiry, for up to this point neither has made public his expectations and images of effectiveness. Their joint invention will have to include a reference to the inquiry itself—to the ways in which they might set criteria for and test their own inquiry as it proceeds.

The PM might suggest, for example, that he does have a notion about the way in which the request from Company X ought to be handled. Perhaps Carlos has one as well. The PM will present his initial idea and the thinking that lies behind it. Carlos might agree but add that he will want to be able to take exception and to suggest alternatives without feeling that the PM will interpret this behavior as a wish on Carlos's part to take sole charge of the business. They ought then to settle on a joint approach. Or, if they find that they cannot do so, to get the further information they will need (including, perhaps, the opinions of other staff) and feed that back into their deliberations. They will want, then, to test out an approach to assigning the task, reflecting ahead of time on the data they will consider in its evaluation.

Such an invention will be a new espoused theory of action for assigning projects like the Company X request. It will not be clear at

the outset just how broad the class of "decisions like X" is to be; the generality of the invention remains to be defined. However, there will still be a considerable gap between the invention and behavior required to carry it out. The invention must still be produced.

A great deal more information will be involved in production than was contained in the brief description of the invention. Like all descriptions of designs, the description of this invention does not take account of the theories-in-use that Carlos and the PM will bring to the tasks called "presenting," "taking exception," "settling on a joint approach," "seeking information," and the like. Nor does the description specify the design for interaction—the timing, density, and rhythm of their meetings. Nor does it specify what is to be done when unexpected information is forthcoming or when unexpected consequences ensue.

The invention, in short, is an incomplete program for action which Carlos and the PM must, in the fine-grained structure of their further behavior, complete. In the nature of things, such a program cannot be specified completely ahead of time.

Given the content of their invention, Carlos and the PM will be engaged in testing both a specific approach to the assignment of Company X's request and an approach to the making of further decisions like that one. They will be engaged, then, in learning what to make of the task-assignment process they devise and what to make of the decision process itself. In both senses, they will be involved in evaluation and generalization. Was the method of assignment effective? And on what criteria were they making this judgment? Did things work out as expected? And what were their respective expectations? If there are discrepancies between expectation and outcome, to what do they attribute these discrepancies?

If they discover that they see the events differently, they will need, once again, to produce the data that led to these perceptions, and to state the assumptions on which their different interpretations were derived from the data. Whatever their learning from the process, they will—tacitly or explicitly, privately or publicly—assign a certain generality to the learning. Do they infer from the experience a program to be followed in assigning further requests from Company X; from all such companies; from all clients? Or are only certain features of the method applicable to future instances? With respect to the decision process, the same sorts of questions will also be pertinent.

In effect, Carlos and the PM (and others, depending on the process they adopt) will again be engaged in a process of mapping and diagnosis. This time they will be mapping and diagnosing the new situation that has resulted from their efforts to produce their invention. Again, they will be dependent on one another for the surfacing of attributions and the generation of data. Again, they will be confronted with the shared task of constructing a picture of the situation their several, interactive behaviors have created.

In this scenario we have described a learning cycle, highlighting the phases of discovery, invention production, and generalization which for clarity's sake we have kept separate from one another. (In real-world processes, these phases would blur into one another; discovery and invention, invention and production, production and generalization, would take the form of interactive loops.) The context of this learning cycle is a small-scale interaction between Carlos and the PM, during which each of them learns. But through their joint inquiry, the two individuals serve as agents of organizational learning. They detect and correct errors in the organizational theory-in-use for task assignment. They also reflect on their earlier attempts to learn about problems of task assignment, drawing from that reflection a new approach to the setting and solving of such problems. They function as agents both for first-order organizational learning and for organizational deutero-learning.

MODEL II THEORY-IN-USE

The scenario of behavioral strategies and consequences that we have just presented is informed by a model for *individual* theory-in-use which we call Model II (Argyris and Schon 1974, pp. 85–93).

Briefly, the governing variables or values of Model II are not the opposite of Model I. The governing variables are valid information, free and informed choice, and internal commitment (Fig. 6.1). The behavior required to fulfill these values also is not the opposite of Model I. For example, Model I emphasizes that the individuals be as articulate as they can be about their purposes and simultaneously control the others and the environment in order to assure that their purposes are achieved. Model II does not reject the skill or competence to be articulate and precise about one's purposes. It does reject the unilateral control that usually accompanies advocacy because the typical

Model II

I Governing variables for action	II Action strategies for actor and toward environment	III Consequences on behavioral world	IV Consequences on learning	V Effectiveness
Valid information Free and informed choice Internal commitment to the choice and constant monitoring of the implementations	Design situations or encounters where participants can be origins and experience high personal causation Task is controlled jointly Protection of self is a joint enterprise and oriented toward growth. Bilateral protection of others	Actor experienced as minimally defensive Minimally defensive interpersonal relations and group dynamics Learning-oriented norms High freedom of choice, internal commitment, and risk-taking	Disconfirmable processes Double-loop learning Frequent testing of theories publicly	Increased effectiveness

FIGURE 6.1

purpose of advocacy is to win. Model II couples articulateness and advocacy with an invitation to others to confront one's views; to alter them in order to produce the position that is based on the most complete valid information possible and to which people involved can become internally committed. This means the actor (in Model II) is skilled at inviting double-loop learning on the part of other individuals.

Every significant Model II action is evaluated in terms of the degree to which it helps the individuals involved generate valid and useful information (including relevant feelings), solve the problem in such a way that it remains solved, and do so without reducing the present level of problem-solving effectiveness.

The behavioral strategies of Model II involve sharing power with anyone who has competence and who is relevant to deciding or implementing the action. The definition of the task and the control over the environment are now shared with all the relevant actors. Saving one's own face or that of others is resisted because it is seen as a defensive nonlearning activity. If face-saving actions must be taken, they are planned jointly with the people involved. The exception would be with individuals who can be shown to be vulnerable to such candid and joint solutions to face-saving yet who need to be protected from others (and since it is done unilaterally) from themselves.

Under these conditions individuals will not tend to compete to make decisions for others, to one-up others, or to outshine others for the purposes of self-gratification. Individuals in a Model II world seek to find the most competent people for the decision to be made. They seek to build viable decision-making networks in which the major function of the group is to maximize the contributions of each member so that when a synthesis is developed, the widest possible exploration of views has occurred.

Finally, if new concepts are created under Model II conditions, the meaning given to them by the creator and the inference processes used to develop them are open to scrutiny by those who are expected to use them. Evaluations and attributions are minimized. However, when they are used, they are coupled with the directly observable data which lead to the formation of the evaluation or attribution. Also, the creator feels responsible for presenting the evaluations and attributions in ways that encourage their open and constructive confrontation.

If the governing values and behavioral strategies just outlined are used, then the degree of defensiveness in individuals, within groups, and between and among groups, will tend to decrease. Free choice will tend to increase as will feelings of internal commitment and essentiality.

The consequence of learning should be an emphasis on double-loop learning, by means of which individuals confront the basic assumptions behind the present views of others and invite confrontation of their own basic assumptions, and by which where underlying hypotheses are tested publicly and are made disconfirmable, not self-sealing. Where individuals function as agents of organizational learning, the consequences of Model II should be an enhancement of the conditions for *organizational* double-loop learning, where assumptions and norms central to organizational theory-in-use are surfaced, publicly confronted, tested, and restructured.

UNAWARENESS OF THE INABILITY TO DOUBLE-LOOP LEARN

The description of Model II appears to us and most of our clients to date to be relatively straightforward. What we, and they, had misjudged was the degree to which they were unaware that they did not know how to behave in Model II ways. People have little difficulty in espousing and believing in Model II (indeed, some may rate it with motherhood and apple pie), but they do have enormous difficulties in making it their theory-in-use, *and* they tend to be unaware of this fact. We emphasize the word "and" because combining awareness of and the desire for Model II with the unawareness of the inability to produce it becomes a serious and unsettling prospect for people. This is especially true of adults who have rarely had to face the fact that they cannot discover-invent-produce-generalize double-loop solutions to organizational problems even after they wish to do so.

Elsewhere we have described in detail the reactions that people tend to have (Argyris 1976a, 1976b) under these conditions. Here, we will attempt to illustrate the point by referring to the case of Carlos. If double-loop learning is to occur, (1) the PM would have to feel less defensive, (2) Carlos would feel free to take risks, (3) each person would search for his inconsistencies and encourage the other to confront them, (4) both would be able to state their views in ways that are

disconfirmable, and (5) both would believe that public testing would not be harmful.

Why didn't these conditions exist? The answer is that both people used Model I theories-in-use that were embedded in and nurtured by an O-I learning system. The idea of publicly testing one's views may appear acceptable at the espoused level, but in an O-I learning system such an act could lead to the other taking advantage of the openness and hence winning. The same would be true if they searched for inconsistencies and encouraged the confrontation of their ideas and feelings.

Under these conditions people do not learn the skills required for Model II even though they may espouse them. Since reflection on this incongruity or gap is unlikely under the conditions of Models I and O-I, individuals will not tend to be aware of their inabilities because they have had no reason to reflect on them. Moreover, even if they wished to reflect, it would be highly unlikely that they would find others who would provide them with the kind of feedback that would be required to detect and correct these types of errors. Hence the lack of ability to double-loop learn and the unawareness of this fact.

But if people do not know how to double-loop learn, then moving toward Model II means that they will have to learn to do so. This means that the learning processes which we have described as discovery-invention-production-generalization are much more complex. For example, if individuals do not know how to discover, then they will first have to discover how to discover, invent ways to discover, produce these ways, and learn and generalize. If they do not know how to invent, they will have to discover how to invent, invent how to invent, produce these inventions, and evaluate and generalize. In other words, double-loop learning requires that the learning process of discovery-invention-production-generalization be applied to each step of the larger learning process. Figure 6.2 shows these wheels within wheels.

For example, President A told five presidents (together forming a learning group) that he had serious doubts about his subordinate B's ability eventually to become the president. As A described the difficulties, the five presidents began to infer that A may manage B in ways that coerce B to take less initiative and be more conforming than A says he wants. President A was surprised to learn this, but through

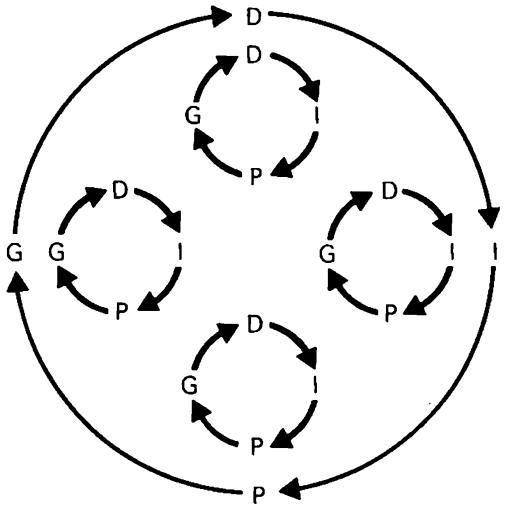

FIGURE 6.2

several hours of discussion, he concluded that he probably was overcontrolling B. The point being made is that A thought that he knew how to discover problems (e.g., B was too passive and dependent), yet the discussion helped him to see that for such problems, he did not know how to discover.

President A then invented a solution. Put simply, it was that if he were overcontrolling, he would become undercontrolling—he would leave B alone. After a lengthy discussion with his peers, A realized that the invention could be counterproductive. For example, B might interpret A's sudden withdrawal differently than A intended. Again, the point is that A soon learned that he did not know how to invent. The same was true for producing and generalizing, not only for A and his group (Argyris 1976a), but also for other groups that we have tested (Argyris 1976b).

MODEL O-II LEARNING SYSTEM

As in the case of Model O-I, we begin the description of Model O-II (Fig. 6.3) with the conditions of error shown in Col. 1. They now interact with Model II theories-in-use (Col. 3) that couple advocacy with inquiry, encourage the surfacing of threatening issues, sanction the making of disconfirmable statements and the testing of these statements publicly.

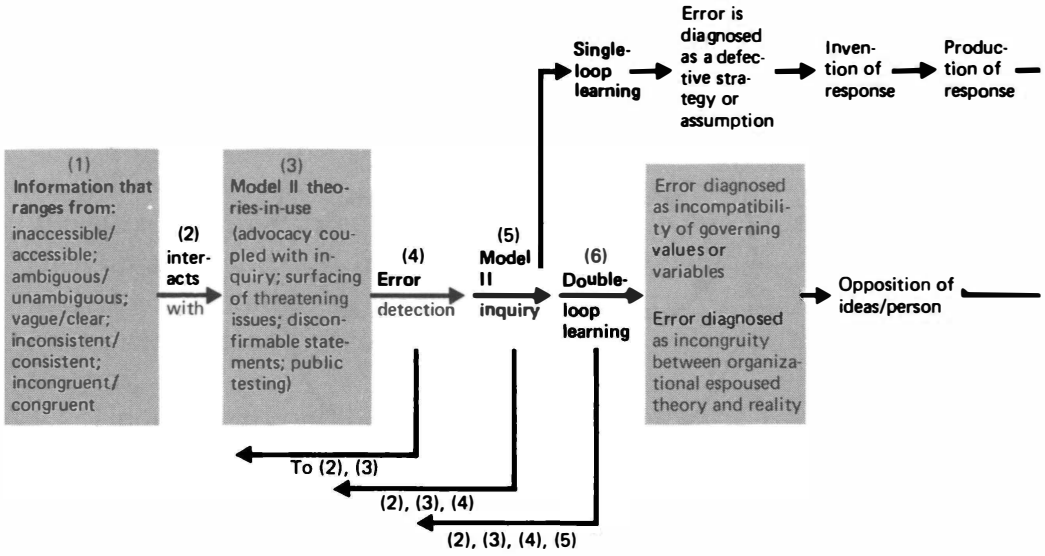

Error is detected and Model II inquiry is begun. Model II inquiry (Col. 5) increases the probability that a corrective response will be made to the conditions for error. As pointed out earlier, this means that mistaken assumptions will tend to be reformulated, incongruities will tend to be specified, ambiguity will tend to be clarified, testability will tend to be substituted for untestability, scattered information will tend to be brought together in concert, information withheld will be surfaced, and information kept impotent for action will be brought into good currency.

The result will be that the conditions of error that were met with dysfunctional responses in Model O-I now will tend to be met with functional responses. Instead of the reaction maintaining or magnifying the errors, errors will now tend to be corrected. The feedback is negative in the sense that it is corrective.

Two kinds of learning are possible in an O-II learning system. The first kind that would be encouraged is single-loop learning. This is relatively straightforward learning because the errors are usually attributable to defective strategies or actions. Consequently, with Model II inquiry it is not too difficult to invent-produce-evaluate effective actions to correct errors. As Fig. 6.3 indicates, inventions are produced to correct the error in strategy or assumption. Since the behavior required fits within the existing theory-in-use, the task of producing it is relatively straightforward. Evaluation then follows: If

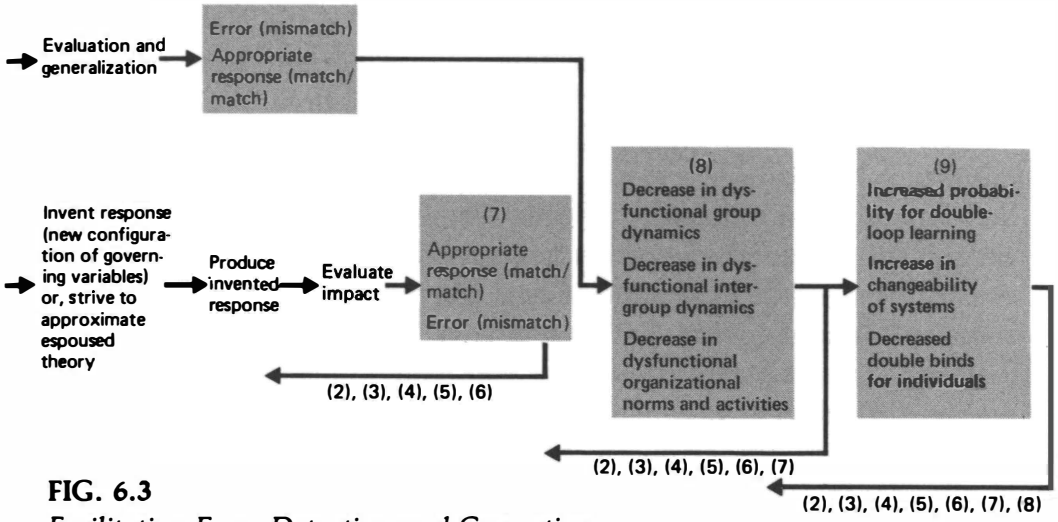

FIG. 6.3
Facilitating Error Detection and Correction

the response corrects the error, the learning is terminated; if the response is a mismatch, the actor returns to diagnosing the error.

The second type of learning is double-loop learning. Here the error is diagnosed as incompatibility of governing values or as incongruity between organizational espoused theory and theory-in-use. The correction of such errors requires the conditions of the good dialectic, which begins with the development of a map that provides a different perspective on the problem (e.g., a different set of governing values or norms). The opposition of ideas and persons then makes it possible to invent responses that approximate the organization's espoused theory. Next, the inventions are produced and evaluated. If the error is corrected, and hence the response is appropriate, the learning cycle ends. If the response is a mismatch there is further inquiry (Col. 7).

Such a learning process should decrease dysfunctional group dynamics because the competitive, win/lose, low trust, low risk-taking processes are replaced by cooperative, inquiry-oriented, high trust and high risk-taking dynamics. Likewise for intergroup relations. Finally, dysfunctional norms and games of deception should decrease, as well as the need for camouflage, camouflage of the camouflage, and the defensive activities described in O-I (Fig. 5.2, Col. 8).

The result should be that the participants will tend to believe that double-loop learning for themselves and their organizations is possible, that organizations are for change, and that the double binds

experienced by individuals should be decreased (Col. 9). Both Col. 8 and 9 reinforce the previous columns and hence we have a learning system that is simultaneously stable and subject to continual change.

One word of warning: As we pointed out in Chapter 3, it is important not to isolate a full cycle of organizational learning from the larger organizational dialectic in which the organization is embedded. In good dialectic, new conditions for error typically emerge as a result of organizational learning, hence the quality of stability combined with continual change. This means that the good dialectic is not a steady state free from conditions for error, but an open-ended process in which cycles of organizational learning create new conditions for error to which members of the organization respond by transforming them so as to set in motion the next phase of inquiry.

TESTS FOR EVALUATING GOOD DIALECTIC

As far as single-loop learning is concerned, we can evaluate good dialectic on a basis of effectiveness. Because norms for performance are given, we can ask whether single-loop learning yields a better or more probable match between expectation and outcome. Similarly, as far as deutro-learning is concerned, we can ask whether there is an increase in learning effectiveness from one learning episode to another. One can make such judgments absolutely, because one shares the organization's norms; or one can make them conditionally, saying, for example, "The organization learned effectively to realize norms of which I disapprove."

In the case of double-loop learning, however, norms are themselves in transition. They cannot be taken as given and used as criteria for learning. It is always possible to say about an organization, for example, "They (agents) learned the right thing," or "They learned the wrong thing." One might even say, "They drew the right (or wrong) lessons from their experience." But such judgments depend on whether or not one shares the values toward which the agents are moving. An outsider makes such a judgment much as a party to the inside conflict might make it, from the point of view of his or her own value stance. There is no Archimedean value-neutral point from which such an evaluation can be made.

Also, as discussed earlier, criteria for organizational inquiry may be defined which apply to each of the principal stages of organizational dialectic. Members of the organization must be able to detect mistakes in organizational theory-in-use, incompatibility in norms, and incongruity between espoused theory and theory-in-use. Hence, assumptions in organizational theory-in-use must be made testable. Members must have access to full and coherent maps of organization, organizational context, and organizational past. In the face of uncertainty, members must be able to restructure their picture of the organizational situation. Because incompatible norms are characteristically surfaced in the form of conflict, members must be able to resolve conflict through advocacy coupled with inquiry.

We can translate these general themes into more specific questions which illustrate features of good organizational dialectic:

- Do members of the organization treat organizational assumptions as testable? And do they search for disconfirming data?
- Are the members of the organization able to integrate, for example, the images of organizational theory-in-use held by employees at different levels and locations with those of management so as to make a single organizational map capable of revealing the interconnections of assumptions and values?
- Do the members of the organization share memories of the organization's past which provide them with a context for the interpretation of present error?

  If not, the organization may continue to respond in single-loop fashion to errors which can yield only to double-loop learning; or, without realizing that they are doing so, they may oscillate for a long time between incompatible values which are horns of a dilemma.
- Has the organization found that its expectations to achieve specified objectives are continually disappointed? If so, there may then be real uncertainty over the proper interpretation of this error. Are the members then able to respond to uncertainty by reflection and by efforts at restructuring their perception of the problem? Are they able to respond, for example, not only by altering work-methods and rates of production, but by reconsidering standards for freedom of work?

- Do the members test for congruence of organizational espoused theory with theory-in-use? Do they test for the compatibility of their norms?
- Do the individual members oppose one another without the awareness that their opposition represents a conflict of organizational values? If so, one side may win without recognizing the costs of victory, and without considering a restructuring of the problem which might allow both sets of values to be met.
- Or do members couple advocacy of their own positions with inquiry into the positions of others? Do they keep open the possibility that conflicting values could be internalized by the several members rather than distributed among them by polarization?

These questions point to features of inquiry which are interconnected. For example, the search for disconfirming data reveals errors which can require the restructuring of organizational assumptions. A coherent map of the organizational present, or the past, can reveal incompatible values which would otherwise be ignored.

THE DIALECTIC CONTRASTED WITH OTHER PERSPECTIVES OF ORGANIZATIONAL RATIONALITY AND CHANGE

The dialectical perspective differs significantly from prevailing schools of thought about organizational change and rationality. It does not agree with those who have written as though it were possible, through organizational development techniques of one kind or another, to reach an organizational stable state in which fundamental conflict or dysfunction would no longer arise. Good dialectic is not a matter of smoothness of operation or elimination of error. On the contrary, its goodness is inherent in the ways in which error is continually interpreted and corrected, incompatibility and incongruity are continually engaged, and conflict is continually confronted and resolved.

The dialectical perspective differs from approaches to organizational rationality which emphasize, as those of many management scientists do, the effective achievement of organizational objectives. From a dialectical perspective, we can recognize that there are zones of organizational experience in which objectives (and their associated

norms) are stable, and where organizational rationality may be understood as the search for effective means. This is what we have called single-loop learning. But the dialectic perspective also focuses attention on incompatibility of norms and objectives which are not resolvable by a search for the most effective means. For norms set the criteria by which effectiveness may be judged.

Welfare economists and decision theorists usually take account of conflicting norms, but they tend to believe that such conflicts may be resolved rationally through trade-off analysis. The dialectical perspective recognizes that some conflicts of norms may be resolvable in this way, but that others may not. These intractable conflicts of norms are organizational dilemmas. Good dialectic entails their resolution through double-loop learning; that is, through organizational inquiry which leads to the restructuring of central elements of organizational theory of action.

Prevailing modes of thinking about organizational rationality tend to assume a framework of stable, compatible objectives for which rational inquiry consists of choosing the most effective means. Or, if such modes do accept the need and the feasibility of trade-off analysis among conflicting objectives, they tend not to recognize that conflicts of objectives may not be resolvable by trade-off analysis, or that such intractable conflicts (if they exist) may still be resolved through organizational inquiry, or that organizational inquiry is likely to lead to new organization/environment situations which give rise to new conflicts of objectives which, in turn, require organizational double-loop learning.

A SCENARIO OF ORGANIZATIONAL DOUBLE-LOOP LEARNING

As mentioned at the outset, we have not been able to find in our experience or to draw from the literature descriptions of O-II learning systems with the degree of concreteness that was possible for Model O-I. Nor can we depend on the reader to fill in our gaps with his or her own knowledge because we predict that few, if any, readers have observed organizations that double-loop learn.

Nevertheless, we thought it might be helpful to describe how such learning might go on in a particular case. In order to set the stage, we must describe an organization with a problem. Let us take, for example, the case of a well-known professional school that discovered

that it was internally fractionated; that its students were not satisfied by their education; that the faculty had doubts about the school as a system; and that both had questions about the long-range viability of the school.

A new Dean was brought into the organization to clear up what most participants acknowledged was a mess. The Dean tried valiantly to produce responses from the faculty which were additive and could form the basis of a new charter for the school, but met with very little success. The faculty began to blame the Dean for the school's inability to change, yet whenever the Dean took initiatives, he was met with stiff resistance from segments of the faculty (depending on the action taken).

Part of the problem was that the faculty held Model I theories-in-use that led to primary and secondary inhibiting loops that prevented people from saying what was in their memories, and/or that led others to say what was in their memories in such a way that it produced defensiveness in others; and/or that led people to polarize issues, take sides, and cancel out the contributions being made. As a result, little progress was made, which acted to confirm in people's minds that the system was beyond change, and led them to place more of the responsibility on the Dean.

To illustrate, we draw heavily from parts of a report that one of us wrote describing important organizational games plus their consequences on the problem-solving activities. One of the most prominent self-protective (and organizationally destructive) devices may be called "distancing." For example, the stance of most of the senior faculty when the Dean first arrived was to distance themselves from any responsibility about the issues of planning and redesign. When the Dean asked for help and cooperation, the most frequent response he received was, "The task of redesigning the School is a Deanly decision." The Dean struggled several times to make such decisions at least a function of the senior faculty, but when he called meetings for that purpose the senior faculty operated beautifully to cancel each other out. They used such devices as polarizing issues, making assertions in ways that made the issues untestable, one-upping each other (for example, "The real problem is . . .," implying that the previous speaker was handling the issues at a skin-surface level), and so on.

What option did the Dean have? One was to confront the problem-solving dynamics and refuse to be caught up in them. The Dean

believed that this alternative was dangerous and could lead to further difficulties. (Our analysis will confirm that fear.) He was soon to find himself in a double bind—damned if he took action, and damned if he did not.

The Dean finally chose the alternative of distancing himself from the faculty. This strategy fitted neatly into the living system. The faculty had a Dean who naturally defended himself the way they defended themselves; he too was now practicing the art of distancing. They could all live with each other now, but the solution would help to assure ineffective problem solving and planning in the future. Again the paradox: A helpful solution was also counterproductive.

There were several other very important consequences. The Dean had to find a group with whom he could talk candidly. He understandably selected the administrators that he had appointed plus a few junior faculty whom he had sponsored. This group became the "inner circle." But the creation of an inner circle divided the school into those who were in and those who were out. Since most of the faculty felt in the latter category, most of them also felt not valued. Being outsiders and feeling not valued gave the faculty structural reasons why distancing was no longer their personal responsibility, as well as rational reasons why they should continue their noninvolvement in such issues as planning. It also helped to produce an interesting reaction to the failure of the Dean's attempts to produce a plan for the school. On the one hand, there was genuine sorrow and dismay on the part of the faculty because they wanted to see the school progress. On the other hand, there was a sense of satisfaction because the failure of the plans "proved" to the faculty that the inner circle could not produce an effective plan, hence assuring that some day they would be needed.

The reader may wonder if the creation of an inner circle could not also lead to the administration becoming surrounded by yes-persons who could provide the top with information that was "shaped" and "managed," if for no other reason than out of concern for the already overpressured and overloaded Dean. Apparently this consequence has not tended to happen to the degree that it could. The people within the inner circle try hard not to distort reality even though it means that they must say things that may be difficult to communicate. But, neither is the administration trying to create the condition of "groupthink." It appears that the Dean's reaction is one of pride in the vitality of his group, but he wishes that they could come up with a structural

solution that would not require that the organization's learning system be confronted.

Again a paradox: The members of the inner circle work hard to minimize the probability of group-think, which leads them to make life even more difficult for the Dean; and he, given his instinct, values their honesty even though it threatens his inner peace. The members are in a bind because either way they behave will make it difficult for the Dean. The Dean is in a double-bind because either way he reacts, he will create difficulties for himself.

During the past several years, the Dean has not been able to produce a plan that the faculty, administration, and other relevant people will accept and to which all would become committed. As a result, the school has continued on its path with little change in its intellectual thrust and structural incoherence. The Dean kept insisting there is a genius in the present organization which few insiders or outsiders could see. Perhaps he needed to see the coherence; otherwise, he would have to do the confronting of people, programs, and practices that were distasteful to all concerned, including himself.

According to this analysis, the Dean's assertion that the organization contains a hidden logic is valid in the sense that the organization's learning system is a creative response to the enormous multilayered defenses. However, his assertion is incomplete because the multilayered defenses prevent a cohesive system from developing that can double-loop learn, make decisions, and implement them with continual monitoring. Here we see another example of the paradox that is characteristic of the school: That which is valid is also dysfunctional.

This early period of distancing between the Dean and the faculty had one advantage: It gave the Dean's office time to straighten out the financial and many of the formal administrative problems. The disadvantage of this period was that it served as a continuous confirmation to the faculty that the school's learning system was brittle and unchangeable. That led them to continue their distancing, which led the Dean to continue his, which made it easier for the faculty to take pot shots at the Dean and vice versa. The faculty probably "knew" that time was on their side. If they waited long enough, the university administration would have to enter the picture.

And this is precisely what happened, because even though the waste and inefficiencies were reduced, the costs rose. The school some day would have to seek new money and/or cut its costs even further.

Both activities would require the kind of organizational learning and problem solving that was difficult to produce in this school.

Enter the President. He told the faculty that the destiny of the school was in their hands, and that to date the faculty was doing a questionable job of planning its future. The faculty responded that the school was financially disadvantaged and that the President did not understand how difficult it was to redesign the educational programs in their profession. No meetings were held by the faculty that would provide data to back up their assertions that the educational planning in their profession was uniquely difficult.

The President, in turn, showed little awareness of the complexities and defensiveness of the school's learning system. He appeared to act as if valid planning could occur if only the faculty wanted to do it and if the Dean were willing to take the initiative. Again, both of these conditions were valid, but the learning system made both highly unlikely.

The fact that the President showed little awareness does not mean that he was not aware. He knew about the school's difficulties, if from no other source than the letters received from the faculty in response to his request for their diagnosis of the school's problems. The probability is quite high, therefore, that the President was withholding this information because, like the Dean and the faculty, he could not see how bringing it up would help rational problem solving. But by not making these issues discussable, and by reinforcing these taboo areas, he (and all of the faculty) made truly effective and innovative planning even more difficult. And again, a valid course of action becomes dysfunctional.

The map presented to date is highly incomplete. Incidents could be presented which, if elaborated upon, would make the map more complex and the analysis more alarming. For example:

- Some senior faculty members voted against their will for appointments because the senior member of the involved field made the demand.
- Some senior faculty voted for a major senior appointment because they knew the individual would not accept it, while others refrained from making their abstention explicit.
- The Dean's office made decisions that had tenure implications

without discussing them with the senior faculty, and the faculty who were especially involved kept silent because they didn't want to lose their credits with the Dean.

- The Dean appointed a committee to create a policy that part-time study at the doctoral level was possible, partially to make legitimate a massive violation of the present policies.

Conditions such as these had an impact on the intellectual climate of the school. Most of the doctoral students worked part time in outside activities which made it difficult for them to do first-class work. Data indicate, for example, that the average of the grades of the students from the school in a large class of 100 students was one full grade below those of the other faculties, yet the students evaluated the grading system and the graders as being very fair (over 85 percent).

Interviews with the best students confirmed a lack of intellectual discussion and confrontation among the students, as well as the games students learn to play to keep the academic standards down so that they can hold outside jobs in order to live. These students found themselves in a race between classroom and office, between field work and consulting, and between pressures of term papers and pressures of reports. They left the world of action to reflect and they soon realized that they had probably jumped from the frying pan into the fire.

The students soon found themselves alone and distanced from each other. Like the Dean and the faculty, they too became full-fledged members of the learning system. For those who were lucky enough to attach themselves to a professor who had a grant or a contract, the isolation from each other remained but the pressures were somewhat reduced. Since these were usually the better students, the probability of their interacting with other students or faculty was reduced greatly. It is not difficult to imagine what happens when harried faculty, who protect themselves by "distancing," encounter harried students who do the same.

The faculty may seek a few dedicated students and work with them closely, leaving the others alone. After awhile, this appeals to those students who had returned to the school in order to do some thinking and reflecting about their practice. Many of these students were genuinely ambivalent about academic life and practice. They wanted both but they tired easily as soon as both academics and prac-

tice came in great amounts. Soon, if for no other reason than to live, their part-time work took on greater importance than their academic work.

One consequence can be seen in the results of an informal study which showed that an increasing number of students were not completing their doctorates. Another result was the accepted practice of openly asking a faculty member who had judged a thesis as inadequate to get off the committee in order for the student to pass. Still another result was that the faculty member passed a paper which, upon questioning, he concluded was not up to par with undergraduate honors theses that he had read in a fine university.

In Part I, we noted that organizations are in constant transaction with their internal and external environments. The transaction was described as a dialectic process, one that unfolded in a sequence of stages where each organizational response to perceived problems contributed to the reaction of a new organizational situation which was in its own way problematic. The examples we gave at that time to illustrate the dialectic began with changes in the environment, the creation of which was not discussed. We then pointed out that the dialectic was not necessarily good or bad; it was simply necessary. We defined good dialectic as that in which organizations take cognizance of their dilemmas and resolve them through double-loop learning.

To return to our example: If the school is to reorganize itself to meet the financial crunch, and if it is to do so in ways that reduce or eliminate the self-sealing, counter-productive problem-solving activities, it requires the creation of a good dialectic. The good dialectic, in turn, requires an O-II learning system.

What would the Dean and the faculty do to begin to move toward a good dialectic? What kinds of tasks would they have to perform to experience double-loop learning? (In the next chapter we will discuss what kinds of interventions may help them to achieve these tasks.)

Let us now suppose that the faculty of the school wanted to begin to alter their O-I learning system toward an O-II learning system in order to deal with the first-order problems like the financial squeeze and the development of a new identity.

In the discussion of the good dialectic, we stated that organizations always exist in an environment. The environment may have varying degrees of impact on the organization. One convenient place

for the faculty to begin would be to explore the nature of their environment and its probable impact on the school. For example, faculty could inquire into the number and types of jobs that will be available in the future, and they could inquire into the financial support available to graduate students plus their attitudes toward that support. It may be, for example, that the major financial support for students is loans. But the more recent students may arrive already burdened with repaying loans for undergraduate or other graduate experiences. Such inquiry may lead to ascertaining how students survive in a high-cost area of the country. One answer may be that they must moonlight. Hence they may be unable to be full-time students even though the regulations require that they be. Such inquiry would naturally lead to inquiry into the admissions policies. For example, at the moment, the students who have been working in their profession for several years are favored over those who have not. But these students are older; they tend to be married and have children, hence they need more money to live. Assuming they are superior students with superior records, then they will be sought after by outside institutions for part-time work. The self-reinforcing cycle is closed.

This suggests another area of inquiry, namely the key organizational norms and policies that may be in conflict. For example, school policy is to require that most students be full-time, yet the students least likely to remain full-time are admitted to the school. Inquiry would also be conducted into the norms and policies about the expected contributions of faculty in teaching and research. For example, school policy is that senior professors are hired to study whatever interests them and to base their courses on these interests. Now, with the financial crunch, it may be necessary to ask senior faculty to teach some of the basic (large) courses. Pressure to get the senior faculty to teach other than their immediate research interests may also come from the students who chose the school primarily because of the reputation of the senior faculty. Also, given the financial crunch and the disappearance of new senior positions, the junior faculty may demand more opportunity to teach research seminars in order to enhance their research and increase the probability that they will find work when they must leave.

Another important conflict of organizational norms is represented by the fact that while the administrators are coerced by the faculty to worry about developing and managing new academic pro-

grams, no new academic programs can be actualized without the vote of the senior faculty. Another source of conflict is the fact that for professors to obtain and hold money for their teaching programs, they must have a political base; i.e., an organizational unit to which the school has committed financial resources. This means that teaching programs must be maintained over a period of time. And this means that the faculty of each program must focus on the differences among them so that the dean's office will have no rational reasons to combine and collapse the programs into fewer units. Each teaching program is therefore in a win/lose situation; if it cooperates or is willing to examine the degree of overlap, it may jeopardize its survival.

These conflicts are related primarily to the instrumental activities of the organization. Each is embedded in, or protected by, a set of norms (interpersonal, group, intergroup, and organizational) that exist at the human level. For example, we described the pervasive defensive strategy of faculty distancing themselves from issues that could upset people. But in order to begin to solve any of the issues just described, personal values and objectives will have to be explored. To illustrate: Some think the policy of admitting older students is supported by faculty who are not research-oriented, when actually the faculty supporting that policy are looking for academic and potential research performance regardless of age. Likewise, some professors maintain that their outside "moonlighting" work may be as valuable as being teaching assistants, while others claim with equal vehemence such work is not valuable because it calls for little reflection and analysis. But what makes these issues especially difficult to explore is that both sides know that the feelings are strong. To surface such feelings would violate Model I governing variables and such Model O-I defensive games as distancing and camouflage.

Another example of multilevel conflicts embedded in each other would be the tactics faculty use to obtain scarce financial resources. Some create and maintain close informal relationships with the Dean. Others who do not have such relationships may try to use influential alumni or student demand as possible levers. What makes it especially difficult to discuss these organizational strategies is that each side makes untested attributions about the motivations of the other. Those who are close to the Dean are seen as apple-polishers. Those who rely on student pressures are seen as using divisive political strategies. And

neither of these perceptions would be explorable in a Model O-I learning system.

Finally, it would be desirable to inquire into the ways in which difficulties with structural organizational arrangements can be reinterpreted as "personality conflicts," thereby increasing the probability that they become undiscussable. For example, some of the so-called personality conflicts may be caused by each faculty member distorting the differences between programs, or negatively evaluating other programs, in order to set the stage to obtain scarce resources. When the others learn of this distortion, they reply in kind. Tensions mount but are suppressed. However, when such feelings erupt in the faculty meetings, even though momentarily, then given the Model I governing variables ("suppress feelings and be rational"), the behavior of those violating the variables is attributed to personality characteristics.

These examples illustrate the requirement that the senior faculty will have to explore thoroughly the degree of congruence between organizational espoused theories and organizational theory-in-use. But we would predict that the moment the faculty attempt to interact in order to problem-solve, they will find that features of the O-I learning system get in their way. For example, they will have to face the Model I tendency to smother possible conflict and emotionality by denying the relevance of important issues, or by polarizing opinion.

Recall that in the cases of Carlos and the others, there was a predisposition to withhold information that was important for understanding and resolving the problems. We have found that this tendency is especially strong with faculty when the information being withheld has a high probability of arousing strong feelings if it is communicated. For example, certain faculty members may have become obsolete; others may be very active but rarely spend adequate time with the students; still others may tend to favor their own graduates for jobs within the school, thereby creating a degree of sameness in intellectual perspective that could be dangerous to students who depend on receiving a broad education about the range of views in their discipline.

It is one thing to point out the areas of need for organizational inquiry (instrumental and interpersonal); it is quite another to implement these recommendations. If the faculty met to discuss these issues,

their Model I theories-in-use, their competitive group and intergroup dynamics, and their organizational games would tend to operate to inhibit learning. Even if the faculty is motivated to double-loop learn, their individual theories-in-use and Model O-I learning system will make it highly unlikely that they would achieve their intentions. For example, the faculty would have to learn to surface their hidden assumptions and the problem areas that they consider taboo. They would have to learn to advocate their positions in ways that encourage confrontation of their views. This, in turn, would require that the faculty learn such skills as speaking in ways that make their views testable (e.g., by providing directly observable data to illustrate their concepts), make their inference processes publicly examinable, and minimize unilateral attributions and evaluations.

These are not, as we shall see, easy skills to develop. Moreover, most of the faculty not only do not have many of these skills, they tend to be unaware that they do not have them. Also, they tend to be unaware that they do not know how to go about learning them.

These and other requirements leading to an O-II learning system may seem formidable, and they are. They tend to produce feelings of frustration in adult participants who may be accustomed to learning new skills quickly. What these participants do not realize is that their major learning skills are for single-loop and not double-loop learning.

Although we have just begun to enumerate the challenges of moving toward an O-II learning system, the reader may justifiably wonder if it is realistic to believe that organizations can make such transformations. There are several reactions to this concern. First, we wonder if organizations have a serious choice of not facing up to double-loop learning, given the pressures on them to improve effectiveness and to confront debilitating conflicts of purpose. Second, we are also impressed with the help that can be given to clients, even though the field is primitive and the technology crude. Finally, as we shall show in the following chapter, we are heartened by the fact that the resistance and concerns individuals may have about transforming the organization's learning system are also sources for energy to learn and to change. The questioning of one's assumptions and the facing of one's resistances to this questioning, are not counterproductive; indeed, they can be the basis for double-loop learning and growth.