

CHAPTER ONE

Introduction*

Organizations pursue intelligence. In that pursuit, they process information, formulate plans and aspirations, interpret environments, generate strategies and decisions, monitor experiences and learn from them, and imitate others as they do the same. Organizations seek to serve their interests and conceptions of self, imposing coherence when they can and exploiting the advantages of confusion, chaos, and ambiguity when the opportunity arises. These efforts are disciplined by the presence of other organizations that, in similar pursuit of their own interests and self-conceptions, provide elements of competition and objects for emulation.

This pursuit of intelligence in organizations is never completely successful. To be sure, organizations sometimes manage to achieve impressive levels of intelligence; but organizational intelligence can also be equivocal, ephemeral, and chimerical. Despite all the paraphernalia of intelligence available to them and powerful incentives to avoid errors, organizations often pursue courses that adversely affect some of the interests represented in an organizational coalition. They routinely make blunders, sometimes huge ones. They often fail to survive very long. They frequently discover that actions they thought were unusually clever were actually preludes to disaster. History is filled with organizational stupidity as well as organizational brilliance.

There are reasons for the failures. The pursuit of organizational intelligence appears quite straightforward, but it is not. The difficulties begin with complications in the concept of intelligence itself. In a general way, intelligence refers to the ability to achieve outcomes that fulfill desires as much as possible. We assess the intelligence of a decision procedure by evaluating the outcomes resulting from its use in terms of organizational desires. An intelligent organization is one that adopts procedures that consistently do well (in the organization's own terms) in the face of constraints imposed by such things as scarce resources and competition.

The assessment of organizational intelligence in such terms is complicated by the fact that desires, actions, and outcomes are all distributed across space and time and connected with each other in intricate ways. The outcomes realized from a particular action are not realized instantaneously but are spread over time. Desires also differ across individuals and groups and over time. They change partly in response to the actions and outcomes that they affect. As a result, any assessment of intelligence depends on the time-and-place perspective chosen. What is intelligent from the point of view of one group's desires over one time period may be quite unintelligent from the perspective of another time and set of desires. And what is intelligent from the point of view of preferences activated when action is taken may be quite unintelligent from the point of view of preferences activated later, particularly after preferences adjust to subsequent experience.

If these conceptual problems were not enough, pursuing organizational intelligence, once defined, is made particularly elusive by three elementary problems: The first problem is the problem of *ignorance*. Not everything is known. The future is uncertain. The consequences of taking one action or another are difficult to anticipate precisely. Sometimes they are difficult to anticipate at all. The future depends in part on the actions of many actors, each of whom is simultaneously trying to anticipate the others. Many of the contingencies can be given probabilistic estimates, but such estimates leave considerable range in the possible futures. Moreover, different people in an organization often anticipate dramatically different futures. The past is also uncertain. It is not uncertain because it still remains to be realized but because it is dimly, inaccurately, or differently recalled. The past is experienced in ways that affect both its interpretation and the memories that are retained about it. History is a story, and the story tellers of the past appear to be as variable as the story tellers of the future. Ignorance about the future and the past is organized around ignorance about the causal structure of the world. Interpretations of why things happen in and to organizations are predominately *post hoc*, socially constructed explanations that seem to shift with changing fashions in interpretation and may provide only weak aids to comprehending a future that will reflect new changes.

The second problem is the problem of *conflict*. Organizations seek intelligence in the name of multiple, nested actors over multiple, nested time periods. The preferences or identities embraced by some participants are inconsistent with the preferences or identities of other participants. What is intelligent action from the point of view of one part of an organization is not intelligent from the point of view of another part. These inconsistencies lead to difficulties in exchanging information and cooperation among members of the organization as individuals act intelligently from their

own points of view. At the same time, preferences and identities endure or change over time, and the actions that are intelligent from the point of view of one time are inconsistent with actions that are intelligent when viewed from a different time perspective. The complications of weighing the claims of different participants and different time perspectives in organizational intelligence are legendary.

The third problem is the problem of *ambiguity* in the evaluative bases of action. Intelligence presumes some accomplishment of what is desired. The desires may take the form of preferences defined over possible outcomes. They may take the form of conceptions of self that are to be fulfilled. The preferences to be pursued or the identities to be enacted are usually assumed to be clear, stable, and exogenous. In organizations, in fact, they are typically neither clear, nor stable, nor exogenous. Preferences and identities are defined ambiguously and measured crudely. They are not stable. They are likely to change between the time action is taken and the time that its outcomes are realized, so that there is a predictable difference between the intelligence of an action *a priori* and its intelligence *ex post*. Moreover, preferences and identities not only change but are likely to change as a result of the process of trying to act intelligently. As a result, the evaluative bases of action – the criteria of intelligence – are ambiguous.

In the face of these problems, organizations adopt various practices and procedures for taking action. In order to arrive at good procedures, they experiment with alternatives, copy the practices of others, consult decision theorists, and hire consultants. They adopt practices that make a credible claim of leading to intelligent actions and defend actions that lead to disasters as reflecting unfortunate outcomes of good decisions. A claim of intelligence for an action, in these terms, is a claim that it resulted from an intelligent procedure. A claim of intelligence for any particular procedure is a claim that the procedure will, in general, lead to favorable outcomes and that there are no systematic errors that are implicit in the procedure itself.

Thus, for example, the claim that rational action is intelligent is a claim that estimating the future consequences of possible current actions and choosing the one with the highest expected value will generally lead to outcomes that are evaluated as desirable in the long run. The credibility of the claim is supported primarily by an indirect route. It can be shown that, in a specific abstractly specified world, a specific rational decision procedure will yield (on average) the best decisions. The demonstrations are in many ways persuasive, but they tend to beg two questions: First, does the abstract world adequately represent the real world? Second, do real organizations when attempting to follow rational procedures in real situations produce equivalently favorable results?

It would be misleading to say that answers to these questions are firmly established by evidence. In general, the idea that intelligence can be achieved by choosing among alternative courses of action on the basis of prior preferences and estimations of likely consequences is widely accepted, but it is also widely questioned. It is not clear that the abstract world of decision theory adequately represents the real world. And when real organizations attempt to follow rational procedures in real situations it is not clear that their actions can be assured to lead to favorable results. The literature is filled with disastrous results of actions based on calculations of expected consequences. There is a substantial body of commentary that notes the difficulties of knowing what alternatives exist, of estimating the future consequences of possible current action, and of establishing valid measures of the values of those consequences. There is also considerable doubt about the way rational action treats future preferences as equivalent to current preferences. Preferences and identities change, and it is future preferences and identities that are relevant to assessing future outcomes.

Claims of intelligence for procedures based on rule-based action and organizational learning are similarly supported by two kinds of indirect observations, rather than by direct evidence. The first observation is that the processes by which rules are established involve more careful assessments of the collective and long-term consequences of individual action than can be expected from individual actors at a particular time. Specific action occurs under pressures of time and is taken by individuals and organizations with incomplete knowledge and under conditions of compressed attention, thus is unlikely to be intelligent. A system of rules allows a separation of the judgments necessary to define intelligent action from the myopic, decentralized, short-term making of decisions by particular actors. The second observation is that rules reflect the accumulation of knowledge in a way that is impossible for individual calculation. They store past experience, experience that extends over more time, individuals, and situations than any one individual's experience can. Rules retain and extend knowledge through differential survival and growth of organizations using more intelligent rules, and they reproduce by spreading from one organization to another.

There is no question that some of the experiences of organizations with learning have been positive. Organizations have been known to improve by learning from experience. They have been known to transfer knowledge to and from others. They have been known to accumulate knowledge and develop procedures for storing and retrieving it, particularly in standard procedures and rules. The evidence for the contribution of learning to intelligence is, however, mixed. Organizations learn and remember the wrong things as well as the right things. They are prone to

superstitious learning. And, sometimes, learning itself produces traps for intelligence.

A persistent theme of the present book is that although they are unquestionably useful, neither rationality nor learning assures favorable outcomes. The point is not simply that outcomes are drawn from a probability distribution thus knowable in advance only up to that distribution. The argument is much more that there are systematic features of rational and learning procedures and systematic features of the interaction between those procedures and key features of organizational life that complicate the realization of intelligence. The pursuit of intelligence is frustrated by irremediable limitations of individuals and organizations, by intrinsic properties of the worlds in which organizations operate, and by unintended traps of adaptive action.

In particular, organizations (like other adaptive systems) are plagued by the difficulty of balancing exploration and exploitation. By exploration is meant such things as search, discovery, novelty, and innovation. It involves variation, risk taking, and experimentation. It commonly leads to disasters but occasionally leads to important new directions and discoveries. By exploitation is meant refinement, routinization, production, and implementation of knowledge. It involves choice, efficiency, selection, and reliability. It usually leads to improvement but often is blind to major redirections.

Both exploration and exploitation are needed for adaptation. Exploration cannot realize its occasional gains without exploitation of discoveries. Exploitation becomes obsolescent without exploration of new directions. Finding a good balance between exploration and exploitation is a recurrent problem of theories of adaptation. In rational choice theories it takes the form of deciding how much should be allocated to search (rather than execution of the currently best alternative). In learning theories it takes the form of deciding how much to experiment with new alternatives (rather than increase competence on old ones). In theories of selection it takes the form of evaluating the rate of generation of new, deviant rules (mutations). Theories of adaptation talk about balancing search and action, balancing efficiency and adaptiveness, balancing variation and selection, balancing change and stability, and balancing diversity and unity.

"Balance" is a nice word, but a cruel concept. Defining an optimum mix of exploration and exploitation is difficult or impossible. It involves trade-offs across time. It also involves trade-offs across people and across levels of a system. The optimum balance may vary from one participant to another. It may vary from an individual to an organization, or from an organization to a population of organizations. We know some things about optimal allocations to exploration and exploitation. For example, in gen-

eral, the shorter the time horizon and the narrower the domain (the lower the level of integration considered), the more the optimum balance shifts toward exploitation. As that example suggests, however, the optimum is not a simple determination but depends on the time perspective and on the part of the system that is the focus of attention.

Even if it were possible to define an optimum mix of exploitation and exploration, achieving that mix would be difficult or impossible. When costs and benefits extend over time, organizations encounter several important biases. Generally, decisions are more localized than their effects, and decision makers tend to focus on nearby effects, rather than distant ones. Similarly, experiential learning is more responsive to effects that are in the temporal and spatial neighborhood than to effects that are more distant. As a result, ordinary rationality and learning are likely to underinvest in exploration. Moreover, adaptation affects subsequent adaptation. As a particular domain is explored, the pool of prospects is likely to be depleted, thus the likelihood of finding a better alternative is likely to decline over time. Competence increases with experience, thus encouraging the status quo. Aspirations change as a result of one's own experience and the experiences of others. Values shift.

Two conspicuous dynamic threats to the exploration/exploitation balance are especially relevant to organizations. The first is the failure trap, a dynamic of excessive exploration. Experiments usually fail. They fail because, on average, new ideas are poor, because incompetence with new ideas makes even good ideas less productive than they might be, and because new ideas are likely to be oversold in order to be adopted, thus to be burdened with excessive expectations. The result is a cycle of failure and exploration, running through a series of new ideas but failing to persist in any one of them long enough to discover whether it might be good.

The second dynamic threat to the balance is the success (or competency) trap – a dynamic of excessive exploitation. The rewards for exploitation tend to be more certain, nearer in time, and nearer in space than the rewards for exploration. Successful experience leads to more experience which leads to greater competence which leads to more success, and so on. The local positive feedback among experience, competence, and success generates traps of distinctive competence and results in inadequate exploration.

The complications of finding an appropriate balance between exploration and exploitation are legendary. Even if they were not, ignorance, conflict, and ambiguity combine with a changing world to make efforts to achieve intelligence endlessly challenging. This book does not resolve these problems, but examines them. It describes some ways of thinking about the relations among organizational situations, actions, and outcomes. The core of understanding the pursuit of organizational intelligence lies in

coming to grips with those relations. In that spirit, the chapters focus on four general topics:

- Chapters 2 to 4 examine ideas about how decisions happen in organizations. They contrast two major conceptions of the bases of action – ideas of choice and ideas of rule-following – and consider the possibilities and limitations of seeking intelligence through such means. They elaborate a modern history of perspectives on action and the various claims made for alternative conceptions of decision making.
- Chapters 5 to 11 examine ideas about organizational learning and other forms of organizational change. They discuss some aspects of the evolution of thinking about social evolution and special features of organizational learning. In particular, these chapters explore a few of the complications involved in specifying and achieving a desirable balance between exploiting what is already known and exploring what might come to be known. They also consider some of the problems involved in comprehending a complex world on the basis of small, biased samples of experience.
- Chapters 12 to 15 examine ideas about risk taking in organizations. They consider the idea of risk preference with its associated notions of risk aversion and risk seeking. They develop a conception of variable risk preferences – risk taking that depends on the relation between aspiration levels and performance – and suggest some reasons why ordinary experiential learning will tend to produce a predilection toward risk aversion for gains and risk seeking for losses. More generally, they develop some conceptions of the role of knowledge, targets, foolishness, and identities in the taking of risks.
- Chapters 16 to 21 discuss some aspects of the giving and taking of advice in organizations. In particular, they consider some complications in consultation and expertise, the dilemmas of wisdom, the limits of cleverness, and the roles of innocence, trust, and senses of self. They examine problems of agency and conflict of interest between advice givers and advice takers. They discuss the limitations and biases of expertise, the possibilities for non-strategic advice giving, and the corruptions of consequentialist thought.

The book straddles a deep canyon that divides studies of organizations. On one side of the canyon are found studies trying to understand how things happen in organizations. Why does an organization do what it does? These studies are variously described as “descriptive,” “positive,” or “behavioral.” They try to report organizations as they actually are. On the other side of the canyon are studies devoted to advising how an organization can be induced to produce better outcomes. How can an

organization be improved? These studies are variously described as “prescriptive,” “normative,” or “instructional.” They try to shape organizational actions to conform to models of proper decisions.

Unless you presume that competitive pressures will eliminate all organizations that fail to pursue optimal courses, the two questions are rather different. The first invites the skills and perspectives of behavioral scientists, trying to characterize the way things really are. What are the major phenomena associated with the generation of organizational action? The second invites the skills and perspectives of decision theorists, operations analysts, and economists, trying to show how things might be improved. How can an organization be made more successful or more durable?

The canyon between these two perspectives is deep. A sacred principle of organization studies is the commandment: Thou shalt not confuse behavioral (or descriptive) assertions with normative (or prescriptive) statements. Yet, many writers on organizations straddle this chasm without notable strain. They describe and they advise. Their descriptions slide into their advice (or vice versa) with the ready lubrication of a common terminology, a mixed audience, and a vigorous market for relevance. If the habit of slipping from the language of observers to the language of reformers without changing the words is an academic sin, then academic hell is crowded with students of organization.

The chapters are written primarily from the point of view of describing organizations. They are written by academics who see themselves as observing and reporting organizational life, for the most part, not reforming it. Nevertheless, many of the chapters are in the spirit of straddling the chasm. They describe how organizations operate, but with an eye to how they might be made to operate more intelligently and another eye to the major problems involved in making any assertion about the existence of a guaranteed route to intelligence. Necessarily, such a spirit invites observations of the many ways in which reforms in the name of intelligence may not serve intelligence, of the many possibilities for things that appear obvious to turn out to be misleading or wrong, and of the complexities in establishing that actions, or even outcomes, are desirable.

At the same time, however, the same spirit invites an effort to say something, however tentative, about how things might be better. We ask how to improve the exploitation of what is known; how human beings of limited capabilities for inference can confront incomplete and biased data on history and still make some sense of their experiences; how organizations can cope with the complexities of experience, particularly the way the intelligence of a particular course of action depends on the courses of action adopted by others, as well as the whole history of actions and responses to them; how it is possible to deal with sparse, redundant, confusing data; how the knowledge and actions of others can inform organ-

izational action; how learning can be retained and spread within an organization.

We ask how organizations can improve the exploration of what might come to be known. We examine the proverbial resistance to change in organizations and conclude that knowledge and experience produce systematic biases against exploration, not because of irrational human rigidity or resistance to change but because most new ideas and practices are likely to be inferior to existing ones, particularly in the short run. In the face of this, we ask how organizations can increase the yield from experimentation with new ideas by improving the average quality of new ideas; how they can increase the quantity of new ideas by stimulating experimentation and the taking of risks; how they can increase persistence with new ideas by buffering new ideas from adverse feedback or from responding to it. The fundamental dilemma remains, however, that the quantity of new ideas and persistence in pursuing them can be increased by changing incentives or by encouraging certain forms of willfulness, ignorance, and fantasy; but increasing the number of good new ideas generally leads at the same time to increasing the number of bad new ideas, and specialists in novelty are generally short lived.

We also ask how greater wisdom can arise from the interaction of advice givers and advice takers in organizations. To a substantial extent, the observations on the giving and taking of advice are built around three themes. The first theme elaborates the idea that understanding is necessarily incomplete, that representing the full complexity of the phenomena associated with organizational intelligence would involve a level of complexity that is inaccessible to human communication. The second theme elaborates the idea that intelligent action requires a mixture of knowledge gleaned from an intimate awareness of the fine detail of a specific context and knowledge gleaned from general analytical thinking and that implementing such a mixture requires a high level of self-discipline on the part of both experts and practitioners. The third theme elaborates the idea that the human spirit requires commitment and autonomy and that great commitment comes not so much from anticipation of great consequences or the pursuit of self-conscious cleverness as from a profound sense of self and the obligations of identity.

If organizational intelligence were easier, the tale would be shorter and life would be simpler. Or, possibly, the tale would be simpler and life would be shorter. In either case, it is comforting that consideration of topics as utilitarian as decision making, learning, risk taking, and the giving of advice lead to such deep puzzles as the fundamental bases of human action, the connections between stability and change, the nature of history, the role of imagination, and the place of poetry in ordinary life. The comfort is, however, more the comfort of Søren Kierkegaard than

that of Jeremy Bentham. If the chapters in this book suggest anything of a general sort, it is that the pursuit of organizational intelligence is an activity in which knowledge can sometimes produce power but more reliably produces humility.

Acknowledgments

* The research has been supported by the Spencer Foundation.