MARCH 1994
M. 23 - 55

theory has long been resisted by students, and generations of economists have struggled to persuade students (and managers) to treat cash outlays and forgone gains as equivalent. The resistance of students has a natural satisficing explanation. Satisficing assumes that people are more concerned with success or failure relative to a target than they are with gradations of either success or failure. If out-of-pocket expenditures are treated as decrements from a current aspiration level (and thus as unacceptable) and forgone gains are not, the former are more likely to be avoided than the latter. A satisficing decision maker is likely to make a distinction between risking the "loss" of something that is not yet "possessed" and risking the loss of something that is already considered a possession.

The tendency to code alternatives as above or below an aspiration level or a status quo has important implications for decision making. Whether a glass is seen as half-empty or half-full depends on how the result is framed by aspiration levels and a decision maker's history. The history is important because aspiration levels—the dividing line between good enough and not good enough—are not stable. In particular, individuals adapt their aspirations (targets) to reflect their experience. Studies of aspiration level adjustment in situations in which information on the performance of others is lacking indicate that decision makers revise aspirations in the direction of past performance but retain a bit more optimism than is justified by that experience. Thus, current aspirations can be approximated by a positive constant plus an exponentially weighted moving average of past experience.

If aspirations adapt to experience, then success contains the seeds of failure, and failure contains the seeds of success. In a very general way, empirical data seem to support such a conception. Although there are some signs that chronically impoverished individuals are less happy than chronically rich individuals, studies of lottery winners reveal that they are no more happy than other people, and studies of paraplegics reveal that they are no less happy than others. This pattern of results has led some people to describe life as a "hedonic treadmill." As individuals adapt their aspirations to their experience, both their satisfactions and their dissatisfactions are short-lived.

The world is more complicated than such a simple model would suggest, of course. Aspirations adapt not only to one's own experience but also to the experience of others. They can become attached not just to the level of reward but to the rate of change of reward. They do not adapt instantaneously, and they appear to adapt upward more rapidly than downward. As a result, deviations in a negative direction seem to be more persistently noticed than positive deviations. This "predisposition to dissatisfaction" is, of course, a strong stimulus for search and change in situations where it exists.

1.3 Theories of Attention and Search

In theories of limited rationality, attention is a scarce resource. The evoked set, of alternatives, consequences, and preferences, and the process that produces the evoked set, take on an importance not found in models of infinitely rational decision makers. Not all alternatives are known, they must be sought; not all consequences are known, they must be investigated; not all preferences are known, they must be explored and evoked. The allocation of attention affects the information available and thus the decision.

Ideas that emphasize the importance of attention are found throughout the social and behavioral sciences. In psychology, the rationing of attention is central to notions of editing, framing, and problem solving "set"; in political science, it is central to the notion of agendas; in sociology, it is central to the notion that many things in life are "taken as given" and serve as constraints rather than as decision alternatives. In economics, theories of search are a central concern of the study of decisions. The study of decision making is, in many ways, the study of search and attention.

1.3.1 The Rationing of Attention

In contrast to traditional societies, which are ordinarily described as short of physical and human resources rather than short of time, the modern world is usually described as stimu-

lus-rich and opportunity-filled. There are more things to do than there is time to do them, more claims on attention than can be met. The importance of scheduling and time, and concerns about "information overload," are distinctive complaints. Industries have arisen to compete for the attention of individuals, as well as to advise people on proper time management. The problems are conspicuously not ameliorated by information technology. Time pressures are further dramatized and probably accentuated by telefaxes, car phones, and systems of electronic mail. Computers seem to have done more to increase information load than to reduce it.

The problems of time, attention, and information management are critical to research on decision making. Limitations on attention and information raise dilemmas for actors in the system and cause difficulties for those who would try to understand decisions. If attention is rationed, decisions can no longer be predicted simply by knowing the features of alternative and desires. Decisions will be affected by the way decision makers attend (or fail to attend) to particular preferences, alternatives, and consequences. They will depend on the ecology of attention: who attends to what, and when. Interested participants may not be present at a given decision because they are somewhere else. Something may be overlooked because something else is being attended to. Decisions happen the way they do, in large part, because of the way attention is allocated, and "timing" and "mobilization" are important issues.

Decision makers appear to simplify the attention problem considerably. For example, they respond to deadlines and the initiatives of others. They organize their attention around well-defined options. Insofar as decisions about investments in attention are made consciously, they are delayed as long as possible. The simplifications do not always seem appropriate to students of decision making. Decision makers are often criticized for poor attention management. They are criticized for dealing with the "wrong" things, or for dealing with the right things at the "wrong" time. Short-run problems often seem to be favored over long-run. Crises seem to preempt planning.

1.3.2 Rational Theories of Information and Attention

Investments in information and attention can be examined using the same rational calculations used to make other investments. No rational decision maker will obtain all possible information (unless it has some direct consumption value—as in the case of rabid sports fans). Rational decision makers can be expected to invest in information up to the point at which the marginal expected cost equals the marginal expected return. The cost of information is the expected return that could be realized by investing elsewhere the resources expended to find and comprehend the current information. There are times when information has no decision value. In particular, from the point of view of decision making, if a piece of information will not affect choice, then it is not worth acquiring or attending to.

Since information is costly, rational decision makers can be expected to look for ways to reduce the average costs of attention, computation, calculation, and search. By assuming that actual decision makers and organizations do in fact make such efforts and are effective in optimizing with respect to information costs, information and transaction cost economists generate a series of predictions about the organization of communication, incentives, contracts, and authority. For example, they consider the possibilities for using other resources to "buy" time. Owners hire managers to act in their interests. Managers delegate responsibility to employees. Since agents may not know the interests of those who delegate to them or may not take those interests fully to heart, the use of agents incurs costs of delegation that are experienced in terms of time as well as money.

As a classic example of rationalizing information and its use, consider the design of optimal information codes. A rational code would be designed to minimize the expected cost of sending messages. People typically tell others to "yell if you're in trouble" rather than to "yell as long as you're okay." Yelling takes energy and so should be conserved. Since "being in trouble" is a less likely state than "being okay," energy expenditure is minimized by associating it with the former state rather than the latter. Similarly, if we assume the early American patriot

Paul Revere was an optimal code designer, then we know that he must have calculated the expected cost of alternative codes in signaling an attack by the British as they moved out of Boston. Under such assumptions, his code of "one if by land, two if by sea" tells us that he thought an attack by land was more likely than an attack by sea (assuming, of course, that he assumed the British would not know about his code).

Organizations use many specially designed codes for recording, retrieving, and communicating information. Accounting systems, human resource management systems, and inventory systems are examples. But the most familiar form of information code is a natural language. Languages and other codes partition continuous worlds into discrete states. Language divides all possible gradations of hues into a relatively small number of colors. Language recognizes a small set of kin relationships (a different set in different cultures) among the many relations that could be labeled. Insofar as a natural language can be imagined to have developed in response to considerations of the costs and benefits of alternative codes, it should make decision relevant distinctions easier to communicate than distinctions that are not relevant to decisions. Where fine gradations in colors are important for decisions, the language will be elaborated to reflect fine gradations. Where color distinctions are unimportant for decisions, they will tend to disappear.

It is not trivial to imagine a process of code development that will optimize a code or language, and it would not be overly surprising to observe suboptimal codes. Decision alternatives are often ambiguous, overlapping, and changing, as are costs and benefits. Decisions require tradeoffs across time and space that are not easy to make. And languages are likely to endure for some time after decision options have changed. Moreover, there are strategic issues involved. If codes distinguish possible actions efficiently from the point of view of a decision maker, they simultaneously provide a guide for the strategic manipulation of that decision maker's choices. Since natural languages have evolved in the face of these complications, one speculation is that some puzzling elements of languages—particular their ambiguities, inconsistencies, and redundancies—are actually

efficient solutions to the many ways in which the world does not match the simplifications of rational models of information.

Rational theories of attention, information, and information structures have become some of the more interesting and important domains of modern economics and decision theory. They have been used to fashion substantial contributions to the practices of accounting, communication, and information management. They have also been used to predict important features of organizational forms and practices. However, there is a kind of peculiarity to all such theories. Determining the optimal information strategy, code, investment, or structure requires complete information about information options, quality, processing, and comprehension requirements. It requires a precise specification of preferences that resolve complicated tradeoffs over time and space. In effect, the problem of limits is "solved" by a solution that presumes the absence of limits. Behavioral students of attention, search, and information have generally pursued a different set of ideas.

1.3.3 Satisficing as a Theory of Attention and Search

Rather than focus on rationalizing attention and information decisions, behavioral students of attention are more likely to build on ideas of satisficing. In its early formulations, satisficing was commonly presented as an alternative decision rule to maximizing. Emphasis was placed on the step function characteristics of the satisficing utility function. Actually, satisficing is less a decision rule than a search rule. It specifies the conditions under which search is triggered or stopped, and it directs search to areas of failure. Search is controlled by a comparison between performance and targets. If performance falls below target, search is increased. If performance achieves its target, search is decreased. As performance rises and falls, search falls and rises, with a resulting feedback to performance.

Thus, satisficing has close relatives in the psychology of decision making. The idea that decision makers focus on targets to organize their search and decision activities is standard. The "elimination by aspects" model of choice assumes that decision

makers do not engage in tradeoffs, they simply consider each criterion sequentially—usually in order of importance—and eliminate alternatives that do not exceed a threshold. The "prospect theory" of choice assumes that decision makers are more risk-averse when returns are expected to be above a target than when they are expected to be below a target.

FAILURE-INDUCED SEARCH

The most important step in a satisficing model of search is the comparison of achievements to targets. Decision makers set aspiration levels for important dimensions—firms for sales and profits, museums for contributions and attendance, colleges for enrollments and placements. Achievements are evaluated with respect to those aspirations. Failure increases search, and success decreases search. In a pure satisficing model, search continues as long as achievement is below the target and ends when the target is exceeded. A natural modification of the pure model would allow search to vary with the discrepancy between achievement and the target, with a decreasing effect as the discrepancy increases.

There are three principal features of satisficing as a theory of search:

1. Search is *thermostatic.* Targets (or goals) are essentially search branch points rather than ways of choosing among alternatives directly. They are equivalent to discrimination nets or thermostats; they begin and end search behavior. As a result, researchers frequently learn more about the real operational goals of decision makers by asking for their search triggers than by asking about their "goals."

2. Targets are considered *sequentially.* A satisficing search process is serial rather than parallel; things are considered one at a time—one target, one alternative, one problem. Since decision makers generally act as though they assume a solution will be found in the neighborhood of a symptom of a problem, the first alternatives they consider tend to be local. If sales fall in Texas, then they look for the problem and the solution in Texas. In this way, order effects become important, and better alternatives are likely to be overlooked if inferior, but acceptable, alternatives are evoked earlier.

3. Search is *active in the face of adversity.* In many ways, standard decision theory is a passive theory. It emphasizes making the best of a world as it exists. Decision theory instructs decision makers to calculate the odds, lay the best bet they can, and await the outcome. Satisficing stimulates a more active effort to change adverse worlds. A satisficing decision maker faced with a host of poor alternatives is likely to try to find better ones by changing problem constraints. A maximizing decision maker is more likely to select the best of the poor lot.

SLACK

Satisficing theories of limited rationality assume two adaptive processes that bring aspirations and performance close to each other. First, aspirations adapt to performance. That is, decision makers learn what they should expect. Second, performance adapts to aspirations by increasing search and decreasing slack in the face of failure, decreasing search and increasing slack when faced with success.

Such theories predict that as long as performance exceeds aspirations, search for new alternatives is modest, slack accumulates, and aspirations increase. When performance falls below aspirations, search is stimulated, slack decreases, and aspirations decrease. Search stops when targets are achieved, and if targets are low enough, not all resources will be effectively used. The resulting cushion of unexploited opportunities and undiscovered economies—the difference between a decision maker's realized achievement and potential achievement—is slack.

Slack includes undiscovered and unexploited technological, marketing, and cost reduction opportunities. It includes undiscovered and unexploited strategies. Variations in search intensity or efficiency result in variations in slack. Since knowledge about opportunities may not be shared generally within an organization, organizational slack resources may be preemptively expropriated by subunits. Some units may not work as hard as

others. Some managers may fly first class or may have more elegant offices and more support staff. Professionals may become "more professional"; engineers may satisfy their love of a beautiful design rather than build the most efficient machine.

Slack has the effect of smoothing performance relative to potential performance. Slack stored in good times becomes a buffer against bad times—a reservoir of potential performance. Thus, variations in realized performance will be smaller than variations in environmental munificence. Because performance is managed in this way, slack conceals capabilities. The level of slack is difficult to determine, and it is hard to estimate what level of performance can be achieved if necessary. Individuals and organizations that appear to be operating close to their capacities frequently are able to make substantial improvements in the face of adversity. The lack of clarity about the level of slack, however, makes slack reduction a highly strategic activity in which each part of an organization (and each individual decision maker) seeks to have the other parts give up their slack first.

Thus, slack is managed. A decision maker may choose to have slack as a hedge against adversity, to smooth fluctuations in profits or resources, or as a buffer against the costs of coordination. Slack may be used to inhibit the upward adjustment of aspirations. Decision makers deliberately reduce performance in order to manage their own expectations about the future. Even more, they do so in order to manage the expectations of others. They restrict their performance in order to avoid overachieving a target and causing the target to rise.

ELABORATING THE SEARCH MODEL

Not all search by decision makers is due to failure. Social systems and organizations may take a deliberate anticipatory approach to search. They may create "search departments" both to solve problems (strategy, planning, research and development) and to find them (quality control, customer complaints). This search tends to be orderly, standardized, and somewhat independent of success or failure.

However, the simple thermostat model of satisficing search captures some important truths. Failure-induced search, the basic idea of the model, is clearly a general phenomenon. Necessity is often the mother of invention, and decision makers threatened with failure often discover ways to cut costs, produce better products, and market them more effectively. Slack serves as a buffer, accumulating in good times and decreasing in bad times. The simple model of search, which involves comparing changing performance with a fixed aspiration, does not capture all that is known about satisficing search, however.

First, *aspirations change* over time, and they change endogenously. They are affected by the past performances of the particular individual or organization and by the past performances of those individuals and organizations perceived as comparable. In general, as performances improve, so do aspirations; as performances decline, so do aspirations.

Adaptive aspirations and have very general effects on organizations. The way they, along with failure-induced search, tend to bring performance and aspirations together has already been noted. When performance exceeds the target, search is reduced, slack is increased, and the target is raised. On average, this tends to reduce performance. When performance is below the target, search is increased, slack is decreased, and the target is lowered. On average, this tends to increase performance.

Thus the process of target adjustment can be seen as a substitute for slack adjustment. If targets adapt rapidly, then slack and search will not adapt rapidly, and vice versa. By virtue of the adaptation of aspirations, subjective definitions of success and failure (which control search behavior and—as will be developed later—both risk taking and learning from experience) depend not only on current performance but also on current aspirations for performance (and thus on a performance history).

Second, search is *success-induced* as well as failure-induced. When the presence of slack relaxes coordination and control pressures, decision makers are free to pursue idiosyncratic, local preferences. They may act opportunistically or imperialistically. If they are members of an organization, they may assert

independence from the organization or may pursue linkages with outside constituents (professional organizations or community interests). These activities are forms of slack search, stimulated by success rather than failure.

Slack search differs in character, as well as timing, from search under adversity. It is less tightly tied to key objectives and less likely to be careful. It involves experiments that are, on average, probably inefficient, particularly in the short run, relative to the primary goals of a decision maker or organization. Most such experiments are probably disadvantageous, but they allow for serendipity, foolishness, and variation. The outcomes of slack search are likely to have a lower mean and higher variance than the outcomes of failure-induced search or institutionalized search. The possibility that such activities find a protective cover in the "waste" of slack plays an important role in an expanded theory of long-run adaptation.

Third, search is *supply-driven* as well as demand-driven. Search is a possible way of describing information acquisition in decision making, but the metaphor has its limits if search is seen as prospecting, seeking alternatives and information that lie passively in the environment. A significant feature of contemporary life is that information is not passive. In some circumstances, a better analogy for information acquisition might be to mating, where information is seeking users even as users are seeking information (for example, in the purchase of equipment). Or the proper analogy might be to hunting, where information is actively eluding information seekers (for example, military secrets) or where information "seekers" are actively eluding information sources (for example, investors and stock salespeople). In general, the market in information is a joint consequence of behavior by the recipient and behavior by the transmitter of the information. It cannot be understood without considering both sides of the transaction.

The general structure of an expanded model of satisficing search is sketched in Figure 1. It displays the close relations among changes in aspirations, changes in slack, and changes in search, the direct and indirect effects of slack on performance, and the exogenous effects of institutionalized search, supply-side search, and the performance of others on the dynamics of the system.

Figure 1
Expanded Model of Satisficing Search

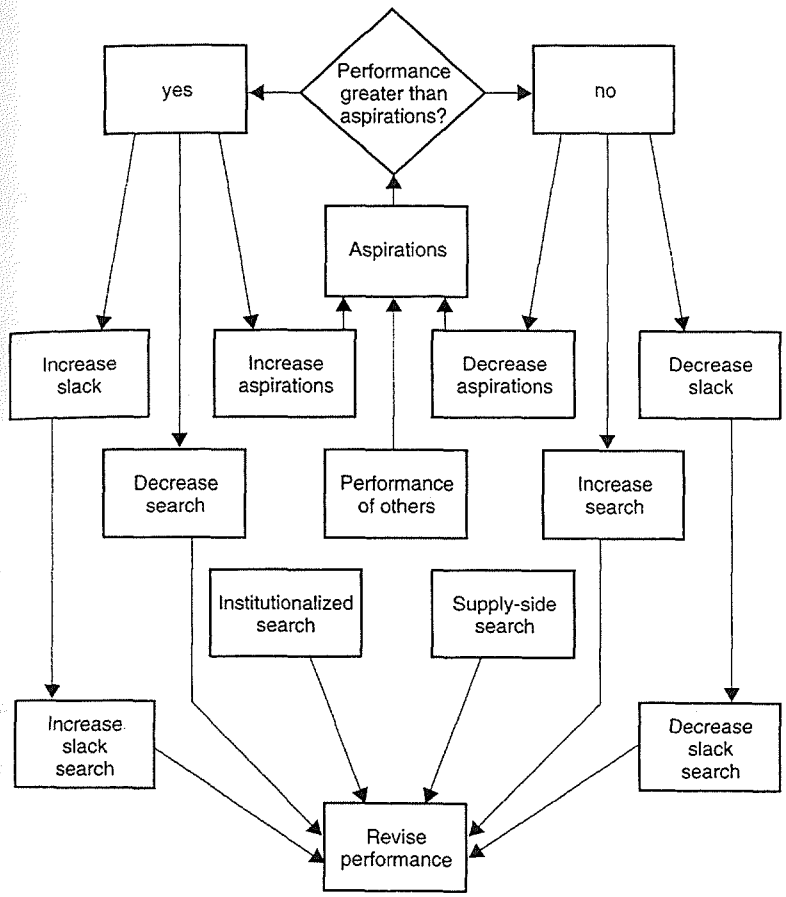

UNDERSTANDING INNOVATION

It is possible to use the general ideas of satisficing search to speculate about the long-run dynamics of individual and institutional change: Do those who have been successful in the past

continue to be successful, or does success sow the seeds of failure? Do the rich get richer or poorer?

There are no simple answers to such questions. Both success and failure stimulate mechanisms that encourage subsequent success, and both success and failure stimulate other mechanisms that encourage subsequent failure. However, an important part of the answer to the stability of success depends on the richness of the search environment. Failure-induced search increases efficiency and reduces foolishness. Success-induced search introduces more risky alternatives. It tends to produce more distant search and introduces bigger changes with lower odds of success. The rich get richer if success-induced search (slack search) gives better returns than failure-induced search or if prior success was produced by either institutionalized search or supply-side search that continues.

In technologically mature worlds, success will tend to breed failure. Slack will produce inefficiencies and unproductive success-induced search. In technologically young worlds, on the other hand, success will tend to breed success. The specific innovation that will provide a breakthrough is hard to identify in advance, so there is a good deal of chance in the outcome from any particular innovation. But slack search provides the resources for relatively frequent experiments, thus increases the chance of an important discovery.

Will there then be persistent innovators? Assuming that all actors are competent, within the satisficing search theory major successful innovations are produced by foolishness, which in turn is produced by a combination of slack (thus success) and luck. Individuals or organizations must be foolish enough to look and lucky enough to find something. A few innovative ideas will be successful, thus marking the individuals and organizations involved as "innovative." Success will lead to slack and thus to more foolish innovative ideas.

As a result, persistently successful organizations will tend to be more innovative than others. However, since most innovative ideas will not be successful, most innovators will not repeat their successes, and their resources will fall, leading them to produce fewer and fewer potentially innovative ideas. Thus,

success in innovation increases the amount of innovative activity. By increasing the amount of innovative activity, it increases the likelihood of new success. But unless the pool of opportunities is rich, it may not increase the likelihood enough to pay the increased costs incurred by the search. Under those circumstances, it leads to long-term decline.

1.4 Risk and Risk Taking

As has been suggested above, understanding risk and risk taking is a serious concern of rational theories of choice. In fact, "risk" is sometimes used as a label for the residual variance in a theory of rational choice. The strategy is to assume that risk preference accounts for any deviation in observed behavior from the behavior that would be observed if decision makers had utilities for money that were linear with money and made decisions by maximizing expected monetary value. This strategy has some appeal for many formal theorists of choice and for many students of aggregate decision behavior.

Behavioral students of decision making are inclined to take a different route. They try to understand the behavioral processes that lead to taking risks. The emphasis is on understanding individual and organizational risk taking rather than fitting the concept into aggregate predictions. As a result, behavioral students of risk are more interested in characterizing the way variability in possible outcomes affects a choice.*

The factors that affect risk taking in individuals and organizations can conveniently be divided into three sets:

1. *Risk estimation.* Decision makers form estimates of the risk involved in a decision. Those estimates affect the risk actually taken. If the risk is underestimated, decisions will reflect greater risk taking than is intended. If the risk is overestimated, decisions will reflect less risk taking than is intended.
2. *Risk-taking propensity.* Different decision makers seem to

*This section draws from work done jointly with Zur Shapira.

have different propensities to take risk. In some choice theories, decision makers are described as having "preferences" for risk. Observations of risk taking suggest that the term "preferences" may incorrectly imply that individual risk propensities are primarily conscious preferences, whereas they appear to arise only partly through conscious choice.

3. *Structural factors* within which risk taking occurs. Both risk estimation and risk-taking propensity are affected by the context in which they occur. Features of organizing for decisions introduce systematic effects into risk taking.

1.4.1 Estimating Risk

Decision makers seek to form estimates of risk that are both technically and socially valid. Technically valid estimates are those that reflect the true situation faced by the decision maker. Socially valid estimates are those that are shared by others, are stable, and are believed with confidence. Neither technical nor social validity can be assured, nor can either be described as distinct.

IMPROVING TECHNICAL VALIDITY

Decision makers typically attribute uncertainty about outcomes to one or more of three different sources: an inherently unpredictable world, incomplete knowledge about the world, and incomplete complete contracting with strategic actors. Each produces efforts to reduce uncertainty.

Inherently Unpredictable Worlds. Some uncertainties are seen as irreducible, inherent in the mechanisms of the universe. For uncertainties that are thought to arise from inherently uncertain environmental processes, decision makers try to judge the likelihood of events. There are numerous studies of individual estimates of the likelihood of uncertain future events. In general, the studies indicate that experienced decision makers are by no means helpless when it comes to estimating future

probabilities. They do rather well in situations in which they have experience.

On the other hand, the mental machinery they use to anticipate the future contains some flaws. For example, future events are rated as more likely to the extent that similar events can be remembered in the decision maker's own past. This is one of the reasons why experienced decision makers do reasonably well in the domain of their experience. The sample from which they draw is related to the universe about which they make predictions. Biases are produced by differences between the universe of relevant events and the sample stored in memory.

Decision makers also assess the likelihood of an event by considering how closely it conforms to a prototypical image of what such an event would look like. Events are judged to be more likely to the extent they are "representative." The most prototypical events are, however, not always the most frequent. In particular, decision makers tend to overlook important information about the base rates of events. Even though the greatest hitters in history were successful only about 40 percent of the time in their best seasons, there is a tendency to expect great baseball hitters to hit whenever they bat, because hitting is what is prototypical of great hitters. Similarly, although great designers produce exceptional designs only a few times in a lifetime, every failure of a great designer to produce a great design is experienced as a surprise.

There are indications that decision makers, in effect, seek to deny uncertainty by focusing on events that are certain to occur or certain not to occur and by ignoring those that are highly uncertain. This is accentuated by the tendency to round extreme probabilities either to certainty or to impossibility. Very few decision makers have the experience necessary to distinguish an event with a probability of 0.001 from one with a probability of 0.00001, although the difference is extremely large and, in some cases, critical.

Incomplete Knowledge. Decision makers tend to exaggerate their control over their environment, overweighting the impacts of their actions and underweighting the impact of other factors,

including chance. They believe things happen because of their intentions and their skills (or lack of them) more than because of contributions from the environment. This tendency is accentuated by success. As a result, although decision makers certainly recognize that some uncertainties are unresolvable, there is a strong tendency to treat uncertainty as something to be removed rather than estimated.

Some of these "avoidable" uncertainties are seen as a result of ignorance or lack of information, incomplete knowledge of the world. For uncertainties that arise from gaps or ambiguities in their knowledge of the environment, decision makers assume that uncertainty can be removed by diligence and imagination. They try to judge and, if possible, improve the quality of information. They have a strong tendency to want their knowledge about what will happen to be couched in terms that deny doubt. They are more likely to seek to confirm their existing information than to acquire or notice disconfirming information. For example, purchasing agents spend a few minutes forming an impression of a potential product, then devote the rest of their time to seeking information consistent with their initial hypothesis.

Since their strategies for understanding uncertain worlds involve forming firm estimates, decision makers appear to prefer stories to more academic information. They prefer information about specific cases to information about general trends. They prefer vivid information to pallid information. They prefer concrete information to abstract statistics. When confronted with inconsistent information, they tend to rely on one cue and exclude others from consideration.

Incomplete Contracting. Some uncertainties are seen as a result of incomplete contracting, the failure to establish understandings with critical people in the environment. Many of the other actors in the environment have interests at variance with those of any particular decision maker. Each decision maker acts on the basis of the probable actions of the others, knowing that they are doing the same. The resulting indeterminacy leads to intelligence systems designed to spy on the intentions of others. It leads to the pursuit of resources to remove dependence on them. And it leads to negotiations to bind others to

desired future actions, rather than to efforts to predict them probabilistically.

The tendency to negotiate and control the environment rather than predict it is consistent with what has already been observed. Uncertainty is treated the same way any other problem is treated—as something to be removed. Decision makers seek control over the uncontrolled part of their environments. Deadlines and guarantees are more common than time-dependent or performance-dependent variable prices, and the latter are more common than time and performance gambles.

IMPROVING SOCIAL VALIDITY

Individuals, social systems, and systems of knowledge all require reasonable stability and agreement in understandings of the world. Without such social validity, decision makers may have difficulty acting, and social systems may have difficulty enduring. The social robustness of beliefs is threatened by the ambiguities of experience and meaning and by the numerous alternative interpretations of reality that can be sustained. Processes toward differentiation persistently break down tendencies toward agreement. Successes lead to decentralization and experimentation in beliefs; failures lead to rejection of beliefs and disagreement.

Countering these pressures toward heterogeneity and instability are an assortment of mechanisms fostering shared and stable estimates of risk. Experience is edited to remove contradictions. Individuals recall prior beliefs as more consistent with present ones than they are. Incongruous data or predictions are likely to be forgotten. Information is gathered to sustain decisions rather than change them. Beliefs are adjusted to be consistent with actions. They are shaped by the beliefs of others.

Preferences for vivid and detailed information and for redundant, overly idiosyncratic information fit this picture of augmenting robustness and building confidence. Detailed stories tend to be filled with redundant and arguably irrelevant information, thus probably inefficient and misleading from the standpoint of making more valid estimates of risk. Nevertheless, decision makers show a preference for detailed stories. In-

sofar as the goal of the decision process is to see the world with confidence rather than accuracy, the double counting of evidence becomes an asset rather than a liability. In social contexts, this justification could possibly be explained as the confounding of social influence with personal preference, but the same kind of effect seems to occur even within individuals who are merely trying to justify their choices to themselves. Confidence increases with the amount of information processed, even though accuracy typically does not.

The view of decision makers as seekers of stable, shared estimates in which they can have confidence is consistent with research on reactions to alternative gambles. At one point, it was speculated that decision makers might be averse not just to uncertainty about outcomes but also to uncertainty about the probabilities of those outcomes. In fact, people seem to seek not certainty of knowledge but social validity. They actually reject clear bets in favor of those with ill-defined probabilities in domains where they feel their estimates and actions are based on valid beliefs. They avoid bets with ill-defined probabilities in domains where they lack such a sense of socially valid knowledge or competence.

1.4.2 Risk-Taking Propensity

The level of risk taking observed in organizations is affected not only by the estimation of the risk but also by the propensity of a risk taker to seek or avoid a particular level of expected risk. Consider four different understandings of risk-taking propensity: (1) risk-taking propensity as a personality trait, (2) risk-taking propensity as a reaction to targets, (3) risk-taking propensity as a reasoned choice, and (4) risk-taking propensity as an artifact of reliability.

RISK-TAKING PROPENSITY AS TRAIT

In one interpretation of risk-taking propensity, propensities for risk are described as individual traits. For example, in many theories of rational choice, particularly those in which risk is measured by nonlinearities in the utility for money, individuals

are assumed to be risk-averse. They are assumed to prefer an alternative that will yield a given return with certainty to any alternative having the same expected value but some chance of higher and lower returns. The assumption of risk aversion is sometimes taken as an unexplained attribute of human beings, sometimes linked to an assumption of decreasing marginal utility of money, sometimes given a somewhat casual competitive advantage survival interpretation.

If people are risk-averse, it is argued, risk taking must be rewarded. Thus, it is expected that risky gambles will be accepted only if they have higher expected returns than those without risk or, more generally, there should be a positive relation between the amount of risk in an investment and the return provided. The argument is impeccable if one accepts the risk-aversion trait assumption and an assumption that markets in risk are efficient. Such assumptions are not universally accepted, and direct observation often produces a negative correlation between risk and return. The assumptions seem to have somewhat greater merit in narrow finance markets than elsewhere—or at least somewhat greater acceptance.

Skepticism about a generic trait of risk aversion, however, does not preclude the possibility that any one individual has a risk-taking propensity that is stable over time but that propensities vary among individuals. In this interpretation, different individuals have different characteristic tastes for risk, some being inherently more risk-averse and some more risk-seeking. Those tastes for risk are seen as established relatively early in life and to be maintained as stable personality traits in adulthood.

The distribution of risk takers in a population (e.g. in a given organization), therefore, is assumed to be affected primarily by selection. Risk-averse people are assumed to select (and to be selected by) different professions and different organizations from those chosen by people more comfortable with risk. The people who become underwater welders or racing drivers will be different kinds of people from those who become postal workers or professors. Thus the solution to creating an organization with a certain "risk propensity" is to attract the right kind of people.

The evidence for variation among decision makers in individually stable risk-taking propensities is mixed, but it seems plausible to suspect that some such variations exist, that there may be consistent differences among people, even consistent differences among cultures or subcultures. However, the evidence also seems to indicate that, at least within a given culture, the risk-taking effects attributable to trait differences in risk propensity are relatively small when compared with other effects.

RISK-TAKING PROPENSITY AS TARGET-ORIENTED

In most behavioral studies of risk taking, individual risk-taking propensity is not seen as a stable trait of an individual but as varying with the situation. Probably the best established situational effect stems from the way decision makers distinguish between situations of success (or expected success) and situations of failure (or expected failure). Risk-taking propensity varies with the relationship between an individual's position and a target or aspiration level, and thus between contexts of success and failure.

When they are in the neighborhood of a target and confront a choice between two items of equal expected value, decision makers tend to choose the less risky alternative if outcomes involve gains, and the more risky alternative if outcomes involve losses. This is a relatively robust empirical result, true for college students, business executives, racetrack bettors, and small granivorous birds.

When individuals find themselves well above the target, they tend to take greater risks—partly because, presumably, in that position they have little chance of failing, and partly because they may be inattentive to their actions as a result of the large cushion. The risk-taking propensities of decision makers who are well below a target are more complicated, especially when their position puts them in danger of not surviving. On the one hand, as they fall farther and farther below their targets, they tend to take bigger and bigger risks, presumably to increase the chance of achieving their targets. On the other hand, as they come closer and closer to extinction, they tend to become rigid and immobile, repeating previous actions and avoiding risk.

Since falling farther from a target and falling closer to extinction are normally correlated, the effect of failure on risk taking appears to depend on whether decision makers focus attention on their hopes (organized around their aspiration level target) or their fears (organized around their extinction level).

These links between success (outcomes minus aspirations) and risk taking are complicated by two important feedbacks:

> First, outcomes are affected by risk taking. At the least, decision makers who take greater risks realize a higher variance in their returns than those who take lower risks. In situations where risk and return are positively correlated, risk takers will, on average, do better than risk avoiders. In situations where risk and return are negatively correlated, risk avoiders will, on average, do better.
>
> Second, aspiration levels (targets) adapt to outcomes. Success leads to higher aspirations; failure leads to lower aspirations. In general, adaptive aspirations tend to moderate the effects of success and failure by making very successful people less risk taking, and by making unsuccessful people less risk taking. Thus adaptive aspirations smooth system performance and risk taking. Explorations of the dynamic properties and long-run competitive consequences of this system suggest that there are some survival advantages in variable risk preferences when combined with adaptive aspiration levels.

RISK-TAKING PROPENSITY AS CHOICE

In a third view of risk-taking propensity, risky behavior is treated not as a function of personality or of aspirations, but as a reasoned choice. In the spirit of the present chapter, individuals can be imagined as rationally calculating what level of risk they think would serve them best. Consider, for example, risk-taking strategy in a competitive situation where relative position makes a difference. Suppose that someone wishes to finish first, and anything else is irrelevant. Such an individual might want to choose a level of risk that maximizes the chance of finishing first. In general, strategies for maximizing the chance of finishing first are quite different from strategies for maximizing expected value.

For example, suppose one were challenged to a tennis match and given the option of specifying the number of points in the match. Given a choice, how long a game would a rational tennis player choose to play, assuming that the length of the game itself had no intrinsic value? The key to answering this question lies in recognizing how the probability of outscoring an opponent depends both on the probability of winning any particular point and on the length of the game. As the length of the game increases, the better player is more and more likely to win, because the variability in outcomes declines with "sample" size (relatively rapidly, in fact). The game's outcome becomes more and more certain, less and less risky.

Any disadvantaged player (i.e., any player who on average loses, for example, a weaker tennis player or a customer at a casino) increases the chance of reaching a positive outcome by decreasing the number of trials (that is, by increasing the sampling error or risk). That is one reason why better students might prefer majors, courses, and examinations with relatively little random error in their evaluations, and poorer students might prefer majors, courses, and examinations with relatively large random error.

Anticipating somewhat the spirit of Chapter 2, it is also possible to observe that individuals might make a reasoned choice of risk that depends not on calculations of expected consequences but on fulfilling the demands of an identity. A culture might define appropriate risk behavior for different roles. For example, it is sometimes reported that teachers seem to expect (and observe) greater playground risk taking by boys than by girls. Rites of passage into different groups require different risk preferences. Similarly, managerial ideology contains a large number of recommendations about the appropriate levels of risk that should be assumed. Management is often defined in terms of taking risks, acting boldly, making tough choices, and making a difference.

RISK-TAKING PROPENSITY AS AN ARTIFACT OF RELIABILITY

Risks may also be taken without consciousness, as a consequences of unreliability—breakdowns in competence, communication, coordination, trust, responsibility, or structure. Cases of risk taking through lack of reliability are easy to overlook, because they have none of the intentional, willful character of strategic, deliberate, or situational risk-taking. Nevertheless, they can be important parts of the risk-taking story.

For example, risk-taking behavior is influenced by changes in the knowledge of a decision maker. Those effects stem from the relation between knowledge and reliability. Ignorance is a prime source of variability in the distribution of possible outcomes from an action. The greater the ignorance of decision makers or of those implementing the decisions, the greater the variability of the outcome distribution conditional on the choice. That is, the greater the risk. Thus, increases in knowledge have two principal effects on a performance distribution: On the one hand, an increase in knowledge increases the mean performance that can be expected in a decision situation. At the same time, knowledge also increases the reliability of the outcome (that is, decreases the risk in the situation). Thus, as decision makers become more knowledgeable, they improve their average performance and reduce their risk taking.

Similarly, social controls tend to increase reliability, thus decrease risk taking. The mechanisms by which controls grow looser and tighter, or become more or less effective, are only marginally connected to conscious risk taking. In general, reliability increases with education and experience, decreases with organizational size. Organizational slack tends to increase in good times and to reduce reliability; it tends to decrease in poor times and to increase reliability. Diversity in organizational tasks or organizational composition tends to reduce reliability. All of these changes affect the actual level of risk exhibited by decision makers.

1.4.3 Organizational Effects on Risk Taking

Organizations often form the context in which riskiness is estimated and risk-taking propensities are enacted into the taking of risks. That context makes a difference. The forms and practices of organizing shape the determinants of risk and thereby the levels of risk taking observed.

BIASES IN ESTIMATION OF RISK

The estimation of risk by decision makers is systematically biased by the experiences they have in organizations. Decision maker experience is not random but is strongly biased in at least two ways: Decision makers are characteristically successful in their past performance in the organization, and they rarely experience rare events. These two mundane facts produce systematic effects in the estimation of risk.

Success-induced Bias. Organizations provide a context of success and failure, both for individuals and for the organizations as a whole. Success and failure, in turn, affect the estimation of risk. Suppose that all outcomes are a mix of ability and luck (risk). Then biases in the perception of the relative contributions of ability and luck to outcomes will translate into biases in the estimation of risk. Any inclination to overattribute outcomes to luck will be associated with overestimating risk, thus with decreasing risk taking. Similarly, any inclination to overattribute outcomes to ability will be associated with underestimating risk, thus with increasing risk taking.

Research on individual attributions of causality to events indicates that success and failure produce systematic biases in attribution. Individuals are more likely to attribute their successes to ability and their failures to luck than they are to attribute their successes to luck and their failures to ability. They are likely to experience lucky successes as deserved and to experience unlucky failures as manifestations of risk. Persistent failure leads to a tendency to overestimate the amount of risk involved in a situation because of oversampling cases in which luck was bad. Persistent success leads to a tendency to underestimate the amount of risk involved because of oversampling cases in which luck was good.

Since organizations promote successful people to positions of power and authority, rather than unsuccessful ones, it is the biases of success that are particularly relevant to decision making. Success makes executives confident in their ability to handle future events; it leads them to believe strongly in their wisdom and insight. They have difficulty recognizing the role of luck in their achievements. They have confidence in their ability to beat the apparent odds. The same conceits may be found in organizational cultures. Successful organizations build a "can do" attitude that leads people in them to underestimate risk. This "can do" attitude is likely to be especially prevalent in young, successful high-growth organizations where the environment conspires to induce decision makers to believe they know the secrets of success. As a result, successful managers (and others who record their stories) tend to underestimate the risk they have experienced and the risk they currently face, and decision makers who are by intention risk-averse may actually be risk-seeking in behavior.

This organizational inducement of risk underestimation may, of course, be useful for the organization. On the one hand, it is a way of compensating for the negative effects of success and upward aspiration adjustments on risk taking. On the other hand, it is a way of inducing the individually self-sacrificing risk taking that serves the organization and the larger society. In situations where risks must be taken in order to be successful, most of those overconfident decision makers will undoubtedly fall prey to the risks they unwittingly face. But only the overconfident will be heroes. Actors in high-performance, quick-decision, high-risk professions (neurosurgery, air force pilots, investment bankers) all share a professional stereotype of being unusually confident. Overconfidence is still overconfidence and often leads to disaster, but in some situations organizations profit from the individual foolishness that unwarranted self-confidence provides.

Biases in Estimating Extreme Probabilities. As has already been observed, there appears to be a tendency for human subjects to assume that extremely unlikely events will never occur and that extremely likely events will occur. This tendency is accentuated by ordinary experiential learning in an organizational setting. Consider an event of great importance to an organization and very low probability. Individuals in the organization can be expected to estimate the probability of the event and to update their estimates on the basis of their experience.

Suppose, for example, that an event of great importance is so unlikely that it is expected to occur only once every hundred years. Examples might be a disaster in a nuclear power facility, an unprecedented flood, or a dramatic scientific discovery. The rare individual or organization that actually experiences a rare event will come to overestimate the likelihood of the event as a result of that experience. However, most individuals in most organizations will never experience such an unlikely event. As a result, experience will lead most individuals in most organizations to underestimate the likelihood of a very unlikely event.

The effects of this underestimation are twofold. First, in cases where the event being estimated is outside the control of the organization (e.g. natural disasters, revolutions), the underestimation leads to a perversity in planning. The tendency is for plans to ignore extremely unlikely events, to treat them as having no chance of occurring. When planning scenarios exclude extremely unlikely events, they tend to overlook (1) that many of these very unlikely events would have very substantial consequences if they were to occur, and (2) that although each one of these events is extremely unlikely to occur, the chance of *none* of them occurring is effectively zero. Predicting precisely which extremely unlikely event with important consequences will occur is impossible, but some such event will almost certainly occur. Yet plans tend to ignore *all* such events. As a result, plans are developed for a future that is known (with near certainty) to be inaccurate.

Second, in cases where the event being estimated is within the control of an organization, underestimating the likelihood of an extremely unlikely event may have perverse motivational and control consequences. Consider the case of "high-reliability" organizations (e.g. nuclear power plants, air traffic control systems, the space program), where organizations go to great lengths to avoid accidents—to manage the system so that an accident becomes an extremely rare event. In such high-reliability systems, most individual decision makers never experience a failure. They come to think the system is more reliable than it is. This exaggerated confidence in the reliability of the system is likely to lead to relaxation of attention to reliability and to a degradation of reliability over time.

Consider, similarly, research and development organizations looking for a rare discovery. Innovative breakthrough discoveries are extremely unlikely events. Most individuals in research never experience them. They come to think breakthroughs are actually rarer than they are. This reduces the motivation to seek such breakthroughs, and thus further reduces the probability.

Most individuals in these two situations learn over time to modify their estimates of risk in directions that are organizationally perverse. Individuals in high-reliability situations underestimate the danger of breakdown and, as a result, increase the danger. Individuals in breakthrough creativity situations underestimate the possibility of discovery and, as a result, reduce the likelihood. The two situations are not entirely parallel, however. The perversities involved in high-reliability are—at some substantial cost—self-correcting. Degradation of reliability leads to increasing the likelihood that individuals will experience a breakdown and recognize that they have underestimated the danger. On the other hand, the perversities in research are not self-correcting in the same way. Reduced motivation to seek discoveries leads to reduced likelihood of such discoveries, thus confirming the earlier underestimate.

SELECTION ON INDIVIDUAL TRAITS

Insofar as risk-taking propensity is an individual trait, the main way in which organizational risk taking can be affected is by affecting the entrance, exit, and promotion of individuals with particular risk-taking propensities.

Who Enters? Who Leaves? Entry into and exit from an organization are commonly seen as voluntary matchmaking and matchbreaking, acts of deliberate consequential choice. In such a vision, a match is established or continued if (and only if) it is acceptable to both the individual and the organization. Thus, in effect, the match between an individual and an organization continues as long as neither has a better alternative. This hypersimple rational model of entries and exits is, of course, subject to a variety of qualifications of the sort considered in this book. But as long as it is taken as a very loose frame, it may serve to

highlight a few features of the process by which individuals and organizations select each other.

In particular, it is possible to ask whether entry or exit processes are likely to be affected by risk-taking propensity. One possibility is that an organization systematically monitors risk-taking propensity and explicitly includes that consideration in its decisions to hire or retain an individual. If risk-taking propensity is observable, the only question is whether one would expect an organization to prefer risk seekers or risk avoiders. The most common speculation is that organizations, particularly those using formal hiring and firing procedures, tend to prefer risk avoiders to risk seekers. The argument is straightforward: Since big employment mistakes are more visible, more attributable, and more connected to the reward system than big employment triumphs, rational employment agents prefer reliable employees to high-risk ones. The argument is plausible, but very little evidence exists for gauging the extent to which it is true.

A second possibility is that organizations do not (or cannot) monitor risk-taking propensity but monitor other things that are, perhaps unknowingly, correlated with risk-taking propensity. For example, suppose employers seek competence. As they assess competence and secure it, they favor individuals who are able to gain and exhibit competence. Since an important element of competence is reliability—being able to accomplish something within relatively small tolerances for error—competence itself selects individuals by traits of risk-avoidance. Thus, unwittingly, an organization in pursuit of ordinary competence disproportionately selects risk avoiders.

Who Moves Up? If risk taking is considered to be a trait that varies from individual to individual, we need to ask not only which individuals enter or exit an organization but also which individuals move toward the top in a hierarchy. As before, it can be imagined that an organization has some preference for risk-seeking or risk-avoiding managers, monitors the behavior of candidates for promotion, and favors those who have the right traits. Also as before, the most common prediction is that (for

reasons similar to those given above) an organization will tend to favor risk-avoiding managers for promotion. As a result, it is predicted that the average risk-taking propensity of higher-level managers will be less than that of lower-level managers.

Surprisingly enough, the small amount of information available to test the prediction indicates that the prediction is wrong. The average risk-taking propensity of higher-level managers appears to be somewhat higher than that of lower-level managers. One possibility is, of course, that organizations monitor risk-taking propensity and differentially promote managers who are prone to take risks. Alternatively, however, it is possible that risk-prone managers are promoted not because the organization consciously seeks risk-seeking executives but because it promotes those who do particularly well.

To explore how this might come to pass, consider the following simple model: Assume that there is a hierarchy within the organization, that there is competition for promotion, and that promotion is based on comparative reputation. Reputation is accumulated over a series of performances on the job. Each single performance on a job is a draw from a distribution having a mean equal to the individual's ability level and a variance equal to the individual's risk-taking propensity. Individuals accumulate reputations over a series of performances. Their reputations are averages of their realized performances. Whenever a vacancy occurs in the organization, the person with the highest reputation on the next lower level is promoted.

Let us assume that individual risk-taking propensity is a trait (individuals do not consciously choose to take risks, they are simply either risky people or cautious people), and that abilities and risk-taking propensities are independent. Then, as the size of the performance samples becomes very large, the reputations of individuals approach their true abilities. The assignment of individuals to levels is determined entirely by the relative abilities of employees. Average ability increases as you move up the hierarchy, and average risk preference is approximately equal at every level in the organization.

However, in real organizations performance samples are typically rather small. For very small performance samples (with

moderate variability in both ability and risk-taking propensity), reputation no longer depends exclusively on ability but is a joint consequence of ability and risk-taking propensity. If the hierarchy is steep (that is, only a few people are promoted from one level to another), the assignment of individuals to levels is heavily dependent on risk preference. Average ability increases very little as you move up the hierarchy, while average risk preference increases substantially. Thus, a procedure that appears to promote people on the basis of their abilities actually moves them ahead on the basis of the amount of risk they take.

EXPERIENCE, LEARNING, AND RELIABILITY

If experience on a job leads to an accumulation of skills and knowledge, then this cumulative knowledge should both increase average performance and increase reliability, decreasing the variance in the performance. As long as competition, promotion, and order effects are relatively small, people with experience will be more likely to stay in a job and in an organization because of their higher average performance, and the increased reliability associated with longer tenure in a job should be manifested in less risk taking.

Moreover, organizations are adept at cumulating experience across individuals to increase both average performance and reliability. They use rules, procedures, and standard practices to ensure that the experiences of earlier individuals are transferred to newer members of the organization. This process of routinization is a powerful factor in converting collective experience into improved average performance. It is also a powerful influence on reliability and should tend to make the average level of risk taken by individuals within an organization decline as the organization ages.

RISK STRATEGIES

In a competitive world, of course, the positive effects of increases in the mean performance must be weighed against the (potentially negative) performance effects of increased reliability. Increasing both competence and reliability is a good strategy

for getting ahead on average. But finishing first in a large field requires not just doing things others do well but doing something different and being lucky enough to have your particular deviation pay off.

In particular, experience gains that increase reliability substantially and mean performance only a little (e.g. standardization, simplification) are not good for competitive advantage when the number of competitors is large. It may be no accident that while experience (as reflected in years of prior work) and knowledge of standard beliefs (as reflected by success in school) are fair predictors of individual success in organizations on average, very conspicuous success in highly competitive situations is not closely related to either experience or knowledge as conventionally defined.

The competitive situation inside and outside an organization affects optimal risk-taking strategies. Suppose that risk can be chosen deliberately and strategically by individual decision makers competing for hierarchical promotion (as above). Any particular individual's reputation will depend on a sample of performances, and the sample mean will depend on two things: ability (which is fixed) and risk taken (which can be chosen). If a hierarchy is relatively steep and reputations are based on relatively small samples of performances, a low-ability person can win only by taking high risks. But if low-ability persons take high risks, the only way a higher-ability person can win in a highly competitive situation is also by taking substantial risks. If the level of risk can be taken arbitrarily to any level, anyone who wants to get ahead will choose to take maximum risks. In this situation there is no screening on ability at all. The "noise" of risk makes it impossible to detect the "signal" of ability. The average ability level will be approximately the same at all levels in the organization, and the average risk preference at all levels will be identical and high.

It should be observed that fluctuations in the importance of risk taking for hierarchical promotion also have implications for the selection of organizations by individuals. If individuals who are ambitious for promotion can choose organizations based on organizational characteristics, then high-ability indi-

viduals will prefer situations where their ability is correctly identified. They will choose situations where reputation is established through large performance samples, where absolute performance is more important than relative performance, and where strategic risk taking is constrained as much as possible. Thus large, steep hierarchies that use small performance samples to establish reputations will be differentially attractive to low-ability people who are ambitious for promotion.

1.4.4 "Risk Taking" and "Risk Preference"

The concept of "risk-preference," like other concepts of preferences in theories of rational choice, divides students of decision making into two groups. The first group, comprising many formal theorists of choice, treats risk preference as revealed by choices and associates it with deviations from linearity in a revealed utility for money. For this group, "risk" has no necessary connection to any observable behavioral rules followed by decision makers. It is simply a feature of a revealed preference function. The second group, consisting of many behavioral students of choice, emphasize the behavioral processes by which risky choices are made or avoided. This group finds many of the factors in risk taking to be rather remote from any observable "preference" for taking or avoiding risk.

To be sure, decision makers often attend to the relationship between opportunities and dangers, and they are often concerned about the latter; but they seem to be relatively insensitive to probability estimates when thinking about taking risks. Although theories of choice tend to treat gambling as a prototypic situation of decision making under risk, decision makers distinguish between "risk taking" and gambling, saying that while they should take risks, they should never gamble. They react to variability more by trying actively to avoid it or to control it than by treating it as a tradeoff with expected value in making a choice.

Sometimes decision makers take greater risks than they do at other times, but ideas of risk, risk taking, and risk preference are all, to some extent, inventions of students of decision making. Often the taking of risk is inadvertant, as is the avoiding of risk. Decision makers take larger or smaller risks because they make errors in estimating the risks they face, because they feel successful or not, because they are knowledgeable or ignorant, because they find themselves in a particular kind of competition.