COGNITIVE REPAIRS:
HOW ORGANIZATIONAL PRACTICES CAN COMPENSATE FOR INDIVIDUAL SHORTCOMINGS

Chip Heath, Richard P. Larrick, and Joshua Klayman

ABSTRACT

The literature in cognitive psychology has described a variety of shortcomings that prevent individuals from learning effectively. We review this literature and provide examples of a number of organizational practices that may effectively repair the cognitive shortcomings of individuals. We call these practices *cognitive repairs*. We then discuss six tradeoffs that affect the success of cognitive repairs. We close by considering how a cognitive perspective might benefit those who study organizational learning and those who manage it.

Research in Organizational Behavior, Volume 20, pages 1-37.

ISBN: 0-7623-0366-2

INTRODUCTION

In a famous speech, Hamlet declares, "What a piece of work is man. How noble in reason, how infinite in faculties" (*Hamlet*, II, 2). An observer who summarized the psychology of the late twentieth century would probably choose very different phrases to describe the human condition—perhaps, "What fools these mortals be" (*Midsummer Night's Dream*, III, 2).

Are people "infinite in faculties" and "noble in reason"? Herbert Simon won a Nobel prize for arguing that social science must understand the ways that human faculties are *limited*. Instead of being infinite in faculties, Simon's humans could be only "boundedly rational" because their cognitive abilities—their ability to perceive, remember, and process information—were restricted. Well, then, if people are not infinite in faculties, are they "noble in reason"? Cognitive psychologists have spent 30 years examining the actual processes that people use when they collect information, combine it, and draw inferences about their world (Nisbett & Ross, 1980; Kahneman, Slovic, & Tversky, 1982; Holland, Holyoak, Nisbett, & Thagard, 1986). Instead of depicting people as "noble" (or magnificent) in reason, this research has argued that people reason in ways that produce systematic errors. A pessimistic modern Hamlet might combine the observations of these two research streams and describe humans as equipped with primitive hardware and buggy software.

However, outsiders have not always accepted the pessimistic description of human faculties and reason that is found in the research literature. As one skeptic put it, "If we are so stupid, how did we get to the moon?" (Nisbett & Ross, 1980).

How should we resolve the apparent discrepancy between the pessimistic literature on human shortcomings and the optimistic evidence of human accomplishment? One way is to dismiss the laboratory research. Some researchers have argued that the shortcomings that have been documented in the lab are so minor that they do not constitute mistakes of any real consequence (Funder, 1987; Cohen, 1981). Others have argued that individuals are less likely to make errors in natural environments than in contrived laboratory experiments (Anderson, 1991; Gigerenzer, 1996; Cheng & Holyoak, 1985; Hilton, 1995).

We propose another way to resolve the discrepancy. Unlike some researchers, we do not dismiss the examples of limitations, errors, and biases reported in the literature; we assume that individuals are limited, their decision processes are biased, and that they often make costly mistakes on important decisions. We resolve the apparent discrepancy between evidence of individual shortcomings and the empirical fact of moonwalks by observing that *individuals* did not make it to the moon, *NASA did.*

Organizations like NASA may have discovered ways to avoid or repair the individual shortcomings that have been documented by cognitive researchers. Organizations may develop such repairs through deliberate analysis, learn them through trial and error, or discover them through serendipitous accident. In some

cases, repairs may derive from formal academic disciplines like economics or statistics (e.g., Nisbett, 1992; Nisbett, Krantz, Jepson, & Kunda, 1983; Larrick, Morgan, & Nisbett, 1990), but in most cases they will not: They will be ad hoc, intuitive rules that emerge from day-to-day practice. Our thesis, then, is that individuals indeed face cognitive limitations and shortcomings, but that organizations can provide individuals with norms and procedures that mitigate their limitations and reduce their shortcomings.

In this paper we describe a variety of potentially serious shortcomings that have been documented in research on human judgment and reasoning. We focus in particular on learning and hypothesis testing, that is, how people use information to develop and revise their mental model of the world. For each cognitive shortcoming we discuss, we provide examples of organizational practices that may repair this shortcoming. We call these practices *cognitive repairs* to emphasize the fact that they correct some cognitive process that was initially flawed and in need of repair.

We identify potential cognitive repairs to spur researchers to consider how such repairs might look and function. Although at this point, we can make only an anecdotal case for interpreting certain practices as "repairs," we hope that, by pointing out some plausible examples of such repairs, we will prompt researchers in both psychology and organizations to consider more systematically how such repairs might function.

More generally, the concept of organizational cognitive repairs illustrates that researchers may find interesting relationships between individual cognition and organizational practice. These relationships have not received the attention they deserve. On the one side, research in cognitive psychology has largely treated individual learners as "rugged individualists" who face a difficult environment alone, equipped only with their own, flawed cognitive strategies. On the other side, organizational research has largely ignored the literature on individual cognition, focusing instead on issues of motivation or incentives. By studying organizational sources of cognitive repairs, we bring together two frequently disparate literatures and demonstrate how knowledge at one level of analysis can inform the other.

By reviewing individual shortcomings and identifying potential cognitive repairs, we also hope to contribute to the academic and popular literature on organizational learning (Epple, Argote, & Devadas, 1991; Deming, 1982; Senge, 1990; Cohen, 1991; Miner & Mezias, 1996). One important means to facilitate learning *by* organizations is to develop processes that overcome the shortcomings of individuals *within* the organization.

Below, we start with a brief example of the kinds of repairs that we consider in this paper. Then we introduce a framework that describes different stages in the learning process, and we use it to review the literature on individual shortcomings and to suggest potential repairs. As a preliminary reply to Hamlet, we say that even if rugged individuals are unlikely to be infinite in faculties and noble in reason, individuals who have access to organizational and cultural repairs may sometimes appear so.

An Example of Our Approach

Consider one study that might be regarded as an ominous indication of ignoble reasoning by individual experts. Hynes and Vanmarcke (1976) asked seven "internationally known" civil engineers to predict the height of a structure that would cause a foundation to fail; they also asked the engineers to set a 50 percent confidence interval around their prediction so that their confidence interval was wide enough to have a 50 percent chance of enclosing the true failure height. The results were quite sobering: not one engineer correctly predicted the true failure height within his or her confidence interval.

Evidently, the civil engineers thought they knew more than they did—if they had been aware of the limitations of their analysis, they would have set wider confidence intervals and would have predicted the true failure height more correctly. In the psychological literature this kind of finding has been labeled "overconfidence," and it is not an aberration. Similar results have been observed with a number of individual professionals (e.g., Russo & Schoemaker, 1992). In summarizing the evidence, Griffin and Tversky (1992) quipped that experts are "often wrong but rarely in doubt."

To illustrate why this study paints an ominous picture of individual reasoning, consider that (unless you are reading this paper outside) you are sitting in a building that was constructed by civil engineers who were substantially less accomplished than the internationally known experts in the study. Your civil engineers made numerous decisions to ensure the stability and safety of your building; they decided how strong to make its roof supports and how stable to make its foundation. If even expert engineers are overconfident, should you be concerned about your safety?

The answer, we believe, is no. Fortunately, the engineering profession has developed a particular repair, called "safety factors," that mitigate the overconfident reasoning of individual engineers. In an actual assignment civil engineers would precisely calculate the amount and strength of foundation materials necessary to hold a structure of a particular height, then they would multiply their precise answer by a safety factor (i.e., a number between three and eight), and use the larger figure to build the foundation. Were the confidence intervals of the engineers too narrow? Yes. Were they too narrow by a factor of three? No.

Safety factors are an example of the kind of cognitive repair we consider in this paper. An organization (e.g., an engineering firm or the engineering profession at large) provides its members with a repair that helps combat a systematic and potentially serious bias in individual judgment. As a result, the organization shields individuals within the organization from acting on their flawed decisions,

and it shields individuals inside and outside the organization from suffering the consequences.

What is a *Cognitive* Repair?

Organizational repairs can roughly be divided into two classes: (1) motivational repairs increase the energy and enthusiasm with which individuals pursue a task and (2) cognitive repairs improve the mental procedures individuals use to decide which task to pursue and how to pursue it.[1] Organizational research on motivation and incentives can be regarded as the study of motivational repairs (Milgrom & Roberts, 1992; Eisenhardt, 1989). Organizations may need to repair motivational problems in order to encourage individuals to learn (e.g., see Heath, Knez, & Camerer, 1993). For example, individuals may not be willing to experiment with new tasks because they have become endowed with the benefits associated with the old task.

Although previous work has recognized the importance of motivational repairs, it has neglected cognitive repairs. Even when individuals have the right incentives and resources, they may not learn from their experience if they use the wrong mental process to generate hypotheses, collect information, and draw conclusions. The civil engineers who misestimated the stability of the clay embankment were adequately motivated to get the right answer. However, they did not on their own invoke the kind of correctives (e.g., safety factors) that might have made their guesses more appropriately cautious.

REVIEW OF INDIVIDUAL LEARNING AND EXAMPLES OF COGNITIVE REPAIRS

In this section we organize the literature on learning and decision making around three different stages of the learning process. Effective learners must (1) generate hypotheses that explain the causal structure of the world, (2) collect information to distinguish among their hypotheses, (3) draw conclusions that are appropriate and cautious. The boundaries between these stages are fuzzy—they are interrelated and interconnected (Klayman, 1995). However, we distinguish among them because they involve different psychological processes.

Our strategy throughout the review is to consider first the individual then the organization. For each stage of learning, we describe how an ideal individual learner might reason, and review psychological research showing how real individuals depart from this ideal. Then, we describe potential cognitive repairs by which organizations might correct the individual shortcoming in question.

Generating Hypotheses

In the first stage of the learning process individuals must generate hypotheses about the relationships among events. Subject to constraints of time and information, individuals should generate hypotheses that are deep (i.e., by considering causes that are more general or systemic) and broad (i.e., by considering a larger number of potential causes). However, a great deal of psychological research suggests that individuals develop hypotheses that are shallow and narrow.

Individuals Generate Hypotheses that are Shallow Rather than Deep

Individuals Search for Explanations that Make Themselves Look Good

Individuals often conduct shallow searches when they try to explain success or failure because they search in a self-serving way (i.e., in a way that allows them to feel good about themselves). In a meta-analysis of 91 tests of this self-serving bias, Mullen and Riordan (1988) show that individuals typically conclude that their *successes* resulted from stable, internal factors (e.g., ability), but that their *failures* resulted from unstable, environmental factors (e.g., the difficulty of the environment, insufficient effort, or bad luck) (see also Fiske & Taylor, 1991, pp. 78-82).

How might organizations repair self-serving biases? Some repairs may be quite simple: Traders on Wall Street are warned, "Don't confuse brains and a bull market" (Odean, 1996). This compact phrase prompts individual traders to consider the base rate of success in the market, and it makes it more difficult for them to indulge in self-serving explanations for their success.

At Florida Power and Light employees developed a new way to fight self-serving biases after an incident that prominently featured a Japanese inspector for the Deming Prize who later became a folk hero within the company (Walton, 1990, p. 61). To impress the inspector, FP&L managers took him to visit a new facility that had been constructed faster and more economically than any facility in the history of the industry. However, the Deming inspector did not simply accept the results at face value and congratulate them on their "quality" project management; instead, he asked a number of questions to determine *why* they were so successful. The managers' answers were so inadequate that it soon became clear that they did not understand enough about their success to recreate it in the future. The inspector dismissed their "success" in his Japanese-accented English—"you were rucky." Later on his phrase, complete with accent, became a common repair for self-serving interpretations of success.

The Deming inspector deflated a self-serving bias by considering alternative hypotheses for success (e.g., luck rather than skill or knowledge). Traditionally at FP&L, managers were not questioned as long as they achieved good results. After this incident managers were much more likely to be asked to explain their successes. If they could not do so, the verdict would be delivered: "you were rucky" (Walton, 1990, p. 61). The strategic use of the accent was designed to remind managers about the earlier incident where luck produced dramatic results that were unlikely to be repeated.

Individuals Focus on People Rather than Situations

Individuals also generate a shallow set of hypotheses because social settings tend to highlight people as causes. In Western culture individuals typically choose to explain events in terms of people's actions and traits rather than situational factors (Gilbert & Malone, 1995; Ross, 1977; Ross & Nisbett, 1991). In a recent study observers heard another student give a pro-life or pro-choice speech on abortion. Afterward, observers assumed speakers held attitudes consistent with their speeches even though the speeches were derived from scripts written by the experimenters and even though the observers themselves told the speakers which position to speak for (Gilbert & Jones, 1986). Similarly, Deming (1982) describes a company that used a variety of flammable products in their production process. After analyzing the data on fires, Deming found that the fires were a stable and predictable outcome of the production process. However, according to Deming, the company president focused his attentions elsewhere. He "sent a letter to every one of the 10,500 employees of the company to plead with them to set fewer fires" (p. 325).

People's actions are frequently more obvious than their situations. Therefore, when individuals generate hypotheses about why an event occurred, their first hypothesis is likely to be that some person caused it (e.g., Ross & Nisbett, 1991). This tendency to focus on people rather than situations has been documented by so many investigators in so many situations that it has been called the *fundamental* attribution error (Ross, 1977; for recent reviews see Ross & Nisbett, 1991; Gilbert & Malone, 1995).

Organizations might repair the fundamental attribution error by reminding individuals to consider causes other than people, especially the people who are likely to be closest to any problem: front-line workers. For example, an old military adage says, "There are no such things as bad troops, only bad officers" (Cohen & Gooch, 1990, p. 228). Parallel repairs are found in total quality management (TQM). Ishikawa says, "whenever mistakes occur, two-thirds to four-fifths of responsibility rests with management" (Ishikawa, 1985, p. ix). Such maxims may partially repair the fundamental attribution error because they encourage individuals to look beyond the front line. On the other hand, they may simply focus the error on people at a higher level. Thus, a better repair may be one from Deming, who tells managers that of the problems he has seen, "94% belong to the system" (Deming, 1982, p. 315). Ishikawa and Deming both use vivid statistics to overcome the fundamental attribution error even though it is unlikely that either has conducted a precise empirical analysis. Deming's "94%" is particularly noteworthy because of its apparent precision.

Individuals Stop Searching as Soon as They Generate One Hypothesis

Self-serving biases and the fundamental attribution error are special cases of a much broader tendency: Individuals tend to stop searching for a cause as soon as they locate a plausible candidate hypothesis (Gregory, Cialdini, & Carpenter, 1982; Hoch, 1984).

To counter this general tendency, organizations have developed some repairs that are widely applicable across a number of domains. In one technique known as the "Five Whys," workers at Toyota learned to ask "why?" five times before they stopped generating hypotheses. When they did so, they were more likely to find a root cause rather than a superficial one. Imai (1986) illustrates the technique with the following example:

Question 1:	Why did the machine stop?
Answer 1:	Because the fuse blew due to an overload.
Question 2:	Why was there an overload?
Answer 2:	Because the bearing lubrication was inadequate.
Question 3:	Why was the lubrication inadequate?
Answer 3:	Because the lubrication pump was not functioning right.
Question 4:	Why wasn't the lubricating pump working right?
Answer 4:	Because the pump axle wore out.
Question 5:	Why was it worn out?
Answer 5:	Because sludge got in.

Imai argues that by asking "why" five times, workers identified "the real cause and therefore the real solution: attaching a strainer to the lubricating pump. If workers had not gone through such repetitive questions, they might have settled with an intermediate countermeasure, such as replacing the fuse" (Imai, 1986, p. 50). Another illustration of the Five Whys deals directly with the fundamental attribution error: "Problem: He doesn't manage well. (1) Why? He's not on the floor. (2) Why? He's in the equipment room. (3) Why? The newest equipment isn't working. (4) Why? Purchasing gave the supplier a short lead time. (5) Why? Poor planning system" (Forum, 1992, p. 54). In general, when individuals ask "why" the first time, they are likely to develop answers that invoke some salient, recent, or proximal event (e.g., some person's actions). Subsequent whys are likely to cause individuals to think more broadly and situationally.

Although the Five Whys is an admirable cognitive repair because of its power and simplicity, individuals may find it difficult to execute by themselves. When individuals have one good hypothesis in mind, that hypothesis often blocks their ability to see alternatives (Gregory, Cialdini, & Carpenter, 1982; Gnepp & Klayman, 1992; Mynatt, Doherty, & Dragan, 1993). For example, Hoch (1984) found that subjects who generated pro reasons for buying a product had more difficulty generating con reasons immediately afterward.

If individuals find it difficult to generate alternate hypotheses on their own, then organizations may repair shallow search by confronting individuals with others who are expert in asking questions that reveal deep causes. At Microsoft, Bill Gates has by personal example, encouraged a culture that relies on relentless questioning. Says one Windows manager, "you go into the meetings and you come out just sweating because, if there is any flaw, he will land on it immediately and pick it to bits" (Cusumano & Selby, 1995, p. 25). Employees "overuse" terms borrowed from Gates, like "drill down" as a euphemism for "going into more detail" ("What Bill Gates Really Wants," 1995).

A similar cognitive repair is found in the organization that administers the Deming quality prize. Here, official Deming inspectors examine managers using a technique called "single-case bore questions." They begin with broad exploratory queries and then relentlessly delve down into weaknesses and omissions in the answers they receive. Single-case bore questions sometimes identify causes that are quite deep. For example, Florida Power and Light often had to deal with power outages that occurred when a tree fell on a power line and severed it. To improve the reliability of its service, FP&L organized a unit to trim all the trees in sites where damage had occurred, and thus prevent future outages. Managers at FP&L congratulated themselves for creating a procedure that prevented future problems. However, the Deming inspectors were not satisfied with the procedure since it prevented problems only in areas that had already experienced a crisis. They searched for a solution at a deeper level, and asked managers a number of questions about what might be considered forestry! What kind of trees grow in the region? Do palms grow faster or slower than oaks? Managers at FP&L realized they did not know the answers to these questions, and that they had not searched deeply enough to solve their problems. After their experience with single-case bore questions, FP&L managers consulted with foresters and developed a regular maintenance procedure to trim trees based on their growth rates and across the entire region, not just in areas where trees had previously severed lines. After participating in sessions of this kind with the Deming inspectors, managers at the firm learned to ask single-case bore questions in their own internal discussions, thus institutionalizing this cognitive repair (Walton, 1990, pp. 57-63).

Individuals Generate Hypotheses that are Narrow Rather Than Broad

In an ideal world individual learners would not only generate deeper hypotheses; they would also consider a broad rather than narrow set of potential hypotheses. However, even when individuals generate alternative hypotheses, their "alternatives" often differ only slightly from one another, and all lie within the same general frame. For example, participants in one experiment were asked to consider the serious parking problem faced by their university, and they were given time to generate as many solutions as they could (Gettys et al., 1987). Combined, participants generated about 300 solutions that researchers were later able to classify into about seven major categories. One category, for example, suggested ways to reduce demand for parking (e.g., by increasing parking fees) and another suggested ways to use parking more efficiently (e.g., by segregating parking slots according to size). The average participant proposed about 11 solutions but these 11 solutions represented only about three of the seven possible categories. The authors asked an independent panel of experts to compile a complete list of high-quality solutions, and they used this complete list to assess how many solutions were missed by each individual. The typical participant missed from 70 to 80 percent of the high-quality solutions. However, when asked, individuals believed they had missed only 25 percent.

Even experts fail to consider a broad range of alternative hypotheses. For example, one group of researchers showed professional auto mechanics a "fault tree" that listed a number of hypotheses about why a car might not start (e.g., battery, starting system, fuel system, ignition). Some mechanics were presented with a "full tree" that contained seven specific hypotheses, others were given a "pruned tree" that omitted some important hypotheses (e.g. the ignition system). The results indicated that when hypotheses were pruned off the tree, mechanics did not adequately consider them (Fischhoff, Slovic, & Lichtenstein, 1978).

How might organizations repair narrow search by individuals? Individuals might search more broadly if they are cued to think about a problem from different perspectives. At Sharp, employees are told to be "dragonflies but not flatfish." Dragonflies have compound eyes and see things from multiple perspectives at once, but flatfish have large eyes that only look in one direction (Nonaka & Takeuchi, 1995).

The "dragonfly" repair exhorts individuals to consider different perspectives, but this may be difficult for individuals to do by themselves. Organizations might repair narrow search more effectively by encouraging individuals to recruit others who have different perspectives. A good example of this is provided by Bridgestone Tire, which conducts "kokai watches" to generate alternative hypotheses for improving work practices. During a kokai watch a group of up to a dozen people, from different areas of a factory, gather for a few hours to watch others work. In one four-hour watch a dozen people identified 63 potential dangers with a new machine (Walton, 1990, pp. 200-201).

The kokai watch has a number of features that ensure that watchers generate a broad array of hypotheses. First, it mandates a large number of watchers (up to 12). Second, it selects watchers from a variety of different areas—in one kokai watch that examined die and material changes, watchers included a plant visitor, a member of the human resources staff, a chemist, and a project manager. "The idea was that people could observe a process, even those who were strangers to it, with fresh eyes, seeing things that closely involved workers might not" (Walton, 1990, p. 200). Third, it ensures that watchers do not discard hypotheses prematurely. The watchers are instructed to "write down anything, 'Hey, looks like the guy is walking too much,' or 'Looks like he's not handling the knife right'" (Walton, 1990, p. 201). Only after watchers generate hypotheses independently are the results combined and filtered.

Other organizational procedures also repair narrow individual search by ensuring that individuals generate hypotheses independently. For example, when Motorola forms cross-functional teams to evaluate new products, they *do not allow* employees who have participated in one product team to participate in another team with a similar product. This prohibition limits the pool of potential team members in a costly way; evaluation teams involve six to nine people and spend two to three months to develop a business plan for the new product. However, by consciously disregarding previous experience, Motorola allows new teams to develop recommendations independently from previous teams. At the same time, Motorola avoids losing the knowledge of previous "veterans"—they serve as a "review team" that evaluates the recommendations of the newest team.[2] Other repairs ensure that a broad range of alternatives are considered *simultaneously*. Some companies divide a product development team into competing subgroups which develop separate project proposals, and only later recombine to debate the advantages and disadvantages of the independent proposals. Again, this strategy is costly because it is redundant. However, it may have advantages because the built-in independence ensures that different subgroups will approach a problem from different perspectives (Nonaka & Takeuchi, 1995, p. 14).

Collecting Information

In the second stage of the learning process ideal learners collect information to test and revise their hypotheses. There are two main sources of such information: the information that individuals already have in their memory and the information that they collect from the environment. Both kinds of information have potential flaws, but individuals might minimize these flaws if they collected information in a rigorous way. However, learners do not always act as though they are aware of the potential flaws in their information—they frequently collect only a small, biased sample.

Individuals Often Collect Small Samples of Information

Individuals often collect only a limited sample of information because they are constrained by time or attention. In a classic study, Payne (1976) asked his subjects to choose one apartment out of a number of alternatives, each of which was described on several different dimensions (e.g., rent, cleanliness, landlord quality, noise level). When subjects chose among only two apartments, they tended to consider all of the information before they decided. However, individuals searched a smaller and smaller percentage of information as more information became available. For example, one subject, who was deciding among 12 apartments characterized on eight different dimensions, looked at only about 25 percent of the information before making a final choice.

It would be reasonable for individual learners to collect only a small sample of information if they performed a cost/benefit analysis and decided that collecting a large sample was too costly. However, there is evidence that individuals collect only a small sample of information because they systematically underestimate the benefits of larger samples. Tversky and Kahneman (1971) argue that individuals typically believe that small samples will be quite similar to the population from which they are drawn. They labeled this belief the "law of small numbers" to highlight that it contradicts the statistical "law of large numbers," which argues that samples can yield an accurate picture of a population when they are sufficiently large. When individuals believe in the law of small numbers, they assume that any sample will be sufficient, no matter how small.

At the extreme, individuals may not collect any information from the external environment because they believe that they already have adequate information stored in their head. Organizations may overcome this tendency by encouraging or requiring individuals to collect larger samples. This kind of repair is pervasive in writings on TQM. "In promoting statistical quality control, we have used the slogan, 'Let us talk with data'" (Ishikawa, 1985, p. 200). At many TQM companies one of the main principles of the quality effort is "Management by Fact" (Walton, 1990, p. 37).

And TQM not only talks about data, it provides individuals with tools that help them collect and analyze data. For example, six of the "Seven Tools" of TQM provide ways to collect data (e.g., checksheets) or to simplify and display large quantities of data (e.g., histograms, scatter plots, Pareto diagrams, control charts) (Deming, 1982; Imai, 1986; Ishikawa, 1982, 1985; Juran, 1992).

Individuals Collect Biased Samples of Information

Individual learners not only collect small samples of information, they also tend to collect samples that are *biased* (i.e., that are unrepresentative of the larger world). Consider the common claims that "the other line always moves faster" or "it only rains after I wash my car." Unless we want to believe that a malevolent spirit is in charge of such harassment, these examples demonstrate that our memories do not store a random sample of all waiting times or all rainstorms—we are more likely to remember the rainstorms that spoil the finish on our freshly washed car. Even when individuals collect information from the outside world (rather than from memory), they do not always attend to the most relevant and important information. Below, we discuss a number of factors that might lead individual learners to collect biased samples.

Individuals Only Consider Available Information

As indicated by the car wash example, individuals often collect biased samples because they collect information that is easily *available* in memory, for example, because it is especially vivid or recent. The problem is that individuals typically assume that the information that is available is also most frequent, probable, and causally important (Tversky & Kahneman, 1973). This assumption is often wrong. Individuals dramatically overestimate the likelihood of vivid causes of death like accidents or homicides, and they underestimate the likelihood of less vivid causes like disease or strokes. Individuals estimate that accidents caused as many deaths as diseases and that homicides were as common as strokes. In fact, diseases take 16 times more lives than accidents and strokes take 11 times more lives than homicides. Individuals also overweight recent information. They assume that the most recent flood provides an upper bound on possible flood damage, and the purchase of earthquake insurance "increases sharply after a quake and then decreases steadily as memories fade" (Slovic, Fischhoff, & Lichtenstein, 1982, p. 465).

Many organizations repair individuals' tendency to rely on biased, available information by instituting a process that collects information more systematically. At a Motorola division that develops equipment for cellular phone systems, one group realized that an availability bias was causing it to overlook certain customers when it evaluated new products. The unit assigned account managers only to large accounts, so when managers evaluated new products, they primarily considered the needs and requirements of only large customers. However, the unit also served a number of smaller customers that did not have their own account manager. Together, these small customers accounted for a large percentage of revenues. Motorola overcame the availability bias by developing a Feature Prioritization Process; they surveyed customers up to four times a year and then weighted all of the inputs based on customer volume and priority.[3]

Hospitals also have a variety of procedures to force individuals to collect information more systematically. Trauma physicians are often confronted by vivid but potentially misleading information. One doctor states that, contrary to what one might expect, stabbings and bullet wounds are "relatively straightforward affairs" because they leave "clear tracks on the body." Other injuries are more difficult to treat because they leave no visible cues. "It would be all too human to focus on a lacerated scalp—a gory but basically insignificant injury—and miss a fractured thighbone that had invisibly severed a major artery" (Rosenthal, 1994, p. 48). The medical profession has developed a series of strict protocols for trauma situations that allow doctors to quickly collect all the relevant information, not just that which is salient. For example, when a patient first enters the emergency room, physicians follow the "ABCs"; they establish airway, then breathing, then circulation.[4] For situations that are more critical, such as cardiac emergencies, protocols are even more rigorous and specific.

If individuals tend to focus on information that is highly available, it is not terribly surprising that they are frequently unaware of missing information. Even when information is present, learners do not pay as much attention to what *doesn't* happen as what *does* (Agostinelli, Sherman, Fazio, & Hearst, 1986; Newman, Wolff, & Hearst, 1980).

Certain professions and organizations have learned to repair the tendency to ignore missing information. Homicide detectives learn to notice the absence of items at murder scenes, since many murderers take back something that belongs to them after committing the crime. "You look at what's been taken and you find out who it belonged to originally" (Fletcher, 1990, p. 75).

A particularly important form of missing information is the absence of experience with highly unusual events. Bank examiners rarely see a bank fail, nuclear technicians rarely see a meltdown, airline personnel rarely witness a crash (March, Sproull, & Tamuz, 1991; Perrow, 1984). Certain organizations institutionalize procedures that encourage individuals to pay attention to such information despite the fact that such events are unlikely to be available in their own experience. For example, at the Federal Reserve Bank, which certifies the security of banks, senior bank examiners deliberately recount stories of failed banks to keep junior examiners aware that they should be vigilant.[5] At one bank's commercial lending department, senior credit officers would hold seminars and informal brown-bag lunches to discuss past lending mistakes, particularly in areas characterized by unusual or rare events (e.g., "problems with highly leveraged companies, real estate, environmental liability on contaminated property").[6] By forcing individuals to rehearse such information, organizations help individuals learn from vicarious experiences that are rare but highly informative. Furthermore, organizations remind individuals of potentially painful information that self-serving biases would make them prefer to ignore.

Individuals Collect Biased Information Based on Their Preexisting Theories

Research suggests that individuals tend to think of "facts, experiences, and arguments that support a current hypothesis more readily than those that refute it" (Klayman, 1995; see also, Baron, 1988; Kunda, 1990; Nisbett & Ross, 1980). Thus, when individuals collect information from memory, they may focus on information that supports their preexisting theories. Individuals may also do this when they collect information from the external environment. For example, when individuals collect information from others, they often ask specific, directive questions that are likely to elicit the answer they expect (Hodgins & Zuckerman, 1993; Zuckerman, Knee, Hodgins, & Miyake, 1995).

The Chicago Board of Trade has a staff of in-house investigators who scrutinize trades that may violate exchange rules. In these investigations, which are obviously quite sensitive, it is very important that investigators do not collect information that is biased by their initial theories. To repair this tendency, the investigators are trained to avoid questions that can be answered with a yes or no response. "This forces an investigator to ask open-ended questions and allows her to draw out as much information about the situation as possible." By asking open-ended questions, investigators avoid the possibility of directing the interview in a way that elicits only information that is consistent with their preexisting theories.[7]

Some organizations have developed maxims that seem designed to encourage individuals to collect data rather than relying on their (potentially biased) theories. At Bridgestone Tire employees use two Japanese terms: genbutsu (actual product) and genba (actual place) (Walton, 1990, p. 194). These terms remind employees not to rely on their own theories, but to actually go out and investigate the actual product in the actual place where the problems arose. Another group (Forum, 1992) uses a similar cognitive repair they call the three actual rule: (1) Go to the actual place; (2) See the actual problem; (3) Talk to the actual people involved.

Individuals Consider Only Part of the Relevant Information

Finally, individual learners may collect biased samples because they tend to collect information from only one small corner of the universe of information. This arises from basic cognitive processes. Memory is associative—when individuals retrieve one piece of information, they tend to think of other information that is linked to it by strong associations, common features, or similar meaning. Even when individuals collect information from the external environment, they are likely to collect information based on the same kind of associative process. Research in cognitive psychology has shown that individuals attend to and process information more comprehensively when they have a mental *schema* that tells them what information is needed in a given situation and where to find it (Anderson, 1995).

Accordingly, organizations can repair biased information collection by providing individuals with a schema that reminds them of the full range of relevant information. Many schemas of this kind can be found in the financial services industry, where individuals must assess a wide variety of information to determine whether to buy, sell, or lend. At the Federal Reserve Bank of New York, the Bank Examinations group protects the FDIC insurance fund by ensuring that individual banks are in sound financial condition. When reviewing each bank, examiners use a rating system known as CAMEL: they review Capital adequacy, Asset quality, Management, Earnings, and Liquidity.[8] In another bank's commercial loan department, credit analysts use the "Five Cs of Credit": Collateral, Capacity, Capital, Conditions, and Character.[9]

Organizational schemas like CAMEL and the Five Cs are likely to encourage individuals to collect a broader range of information than they would normally collect. It would be very easy for individual learners to collect information on only on the most salient factors (such as cash flow in a loan decision). Although cash flow is certainly important, it can also be misleading or unreliable, particularly in an environment where conditions are changing. By emphasizing the Five Cs, a bank can repair the tendency of individual analysts to neglect information about important variables that are less obvious or are harder to assess. For example, the Five Cs reminds loan officers to consider character—What are the management skills of the owners? Do they have good personal credit records? Although the answers to such questions are quite important, individual analysts might forget to ask them in a numbers-oriented environment like a bank, without a cognitive repair like the Five Cs.

Individuals Who Collect Biased Information Fail to Correct for Bias

We have discussed a number of factors that might lead individual learners to collect biased information. However, even if learners collect biased information, they might still be able to draw effective conclusion as long as they recognized the bias and corrected for it. For example, suppose an individual made the statement, "the other line always moves faster," but then reminded herself that such situations might be overly available in her memory. This kind of correction improves the conclusions drawn from even a biased sample. On the other hand, even if individuals are aware that they have collected biased information, they may not know how to correct for biases after the fact. For example, after individuals ask biased questions and therefore receive biased answers, they do not take into account how much the answers were biased by their initial questions (Zuckerman, Knee, Hodgins, & Miyake, 1995).

Because individuals do not always correct their information for biases, some organizations attempt to ensure that individuals collect unbiased samples from the start. Microsoft requires its software developers to use the same *programs* and *machines* that are used by their customers. For example, programmers who were developing the new Windows NT operating system ran the current day's version of the program as they programmed the next day's version. At Microsoft this process is known as "eating your own dog food." It ensures that developers collect a large, unbiased sample of information about the current state of the program. If Windows NT crashed while a developer was designing a new printer driver, he had to fix the problem with NT before he could return to his driver (Cusumano & Selby, 1995, p. 331). Microsoft also requires developers to use the same machine used by customers, a requirement that has been "controversial at times" because developers like to have the fastest, coolest machines on their desks. However, when developers have better technology than the average customer they collect biased information about how well their software programs perform. One manager said, "every time I've had a project where the developers had hardware that was a generation beyond what customers had, the [software] always had performance problems" (Cusumano & Selby, 1995, p. 347). By requiring developers to use the same machines as their customers, Microsoft forces them to collect an

unbiased sample of information about the operating speed and memory demands of the software they are developing.

Drawing Conclusions

After generating hypotheses and collecting information, ideal learners should evaluate the information they have collected and draw conclusions that are appropriate and cautious. Researchers have suggested three main classes of problems that real individuals face when they interpret evidence. First, they often weigh information in a way that is not statistically appropriate—for example, they emphasize the importance of extreme evidence but they do not emphasize the relative amount of extreme versus non-extreme evidence. A second, even more insidious problem is that individuals use their initial theories to interpret the evidence. While individuals may readily accept information that is consistent with their initial hypothesis, they cast a critical eye on information that contradicts it. Third, as a result of the two previous processes and others, individuals frequently draw conclusions that are overconfident and overly optimistic.

Individuals Weigh Vivid and Extreme Evidence More Heavily

Once individuals have collected information, how should they combine it and weigh it? An ideal learner would weigh information based on the quality of the information. However, actual learners do not always assign appropriate weights to all aspects of the decision. For example, they tend to weigh more vivid, easily imagined information more heavily (Keller & McGill, 1994). They also focus on the extremity or strength of the available information (e.g., the warmth of a recommendation letter) without adequately attending to the amount or weight of the evidence (e.g., the writer's amount of contact with the recommendee) (Griffin & Tversky, 1992).

If individuals tend to overemphasize vivid or extreme information, organizations might prevent this by requiring individuals to consciously classify information according to its appropriate weight. Many companies have internal audit groups that examine the records of company divisions to ensure that they are using proper accounting procedures and spending money on legitimate expenses. An audit usually uncovers a variety of major and minor "exceptions" (i.e., situations where correct procedures were not followed). One auditor says that auditors must be careful not to "place too much emphasis on memorable errors, e.g., an error in the president's expense report or the misuse of the company car." One auditing group repaired this temptation by first classifying each exception as major or minor then consciously ignoring the minor issues.[10]

Consistent with the tendency to overweight the extremity of information and ignore the amount, individuals frequently place higher weight on one vivid case than on a much larger sample of information. Joseph Stalin is reported to have said, "The death of a single Russian soldier is a tragedy. A million deaths is a statistic" (Nisbett & Ross, 1980, p. 43). In a study that supports this observation, Borgida and Nisbett (1977) showed some students a statistical summary of how dozens of students had rated various courses in the previous term. Other students attended a panel discussion, during which two or three upper-division students rated each course on a numerical scale and provided some generic, uninformative comments. Despite the fact that the statistical summary provided students with a larger amount of information, individuals who heard the small sample of vivid information were more likely to change the courses they selected.

Microsoft also discovered that individuals discount large samples of statistical information. At one point, Microsoft started surveying users to see how many of them found it easy to use a particular feature. Software developers often refused to believe the statistics. "The usability group would tell the development group 'Six out of ten couldn't do this.' And the developer's reaction would be, 'Where'd you find six dumb people?'" (Cusumano & Selby, 1995, p. 379). In order to repair this tendency to ignore base rate information, Microsoft made the information more vivid. It built a "usability test lab" where developers can watch real users struggle with new products from behind a one-way mirror. Instead of presenting developers with pallid statistics, the test lab presents them with real people (albeit a much smaller sample). The lab manager says that when developers see a user, "twenty ideas just immediately come to mind. First of all, you immediately empathize with the person. The usual nonsense answer 'Well, they can just look in the manual if they don't know how to use it,' or 'My idea is brilliant; you just found ten stupid people'...that kind of stuff just goes out the door..." (Cusumano & Selby, 1995, p. 379). This cognitive repair is interesting because it uses one kind of bias (overweighting of extreme, or vivid information) to fight another (underweighting of statistical information).

Individuals Use Their Preexisting Theories to Interpret the Evidence

Individuals not only weigh information inappropriately, they also have difficulty interpreting information independently of their preexisting theories. Instead of using the information to test their theories, they use their theories to test their information. This often leads them to discount information that disagrees with their preexisting beliefs.

In a classic demonstration of such discounting, Lord, Ross, and Lepper (1979) selected undergraduates who strongly supported or opposed capital punishment and presented them with two purported academic studies that evaluated capital punishment's effectiveness using very different methods. A study using one method found that capital punishment was effective and a study using the other method found it was ineffective (the researchers counterbalanced which method was associated with which result). Participants applauded the positive aspects of whichever method supported their own preexisting theory, and they critiqued the

"design flaws" in the other. In fact, after receiving mixed results from the two studies, subjects became *more* convinced of the validity of their original position. Seemingly, they regarded the evidence as "one good study that supports my beliefs, and one lousy study that draws the wrong conclusions." Individual subjects thus failed to evaluate the incoming information separately from their preexisting theories. Unfortunately, similar results have been noted with professional scientists (Mahoney, 1976; Koehler, 1993).

One bank helped its loan officers repair the way they interpret evidence by encouraging them to consider a nonstandard theory of lending. In mortgage lending, loan officers often look for reasons to deny loans because loans are difficult to make (they are subject to a mountain of regulations) and potentially quite costly (e.g., foreclosure on a bad loan may cost up to 20% of the property value). Thus, the initial hypothesis in many loan decisions is that an applicant should be denied a loan unless proven otherwise. One mortgage loan department grew at an annual rate of 30 percent by forcing loan officers to consider an alternative to the standard hypothesis. Instead of asking whether an applicant should be *denied* a mortgage loan, it asked whether the applicant should be *approved*. This reversal led the department to develop special programs for qualified applicants who had low incomes or other special circumstances.[11]

Individuals use their theories to develop expectations about what is normal, and they frequently label unexpected events as "problems" or "failures." These labels may be misleading, however, particularly in research and development where unexpected events may point the way to important breakthroughs. One research organization has developed a repair that discourages individuals from thinking that unexpected events are failures (Sapolsky, 1997). Jackson Laboratories breeds mice that exhibit physiological or behavioral traits that are of interest to medical researchers. For example, it sells mice that lack growth hormone to researchers who are interested in understanding the biology of mammalian growth. It found that the animal technicians (e.g., the people who cleaned the cages) often noticed unusual behavior that was scientifically important. The mice that lacked growth hormone were discovered by a technician who noticed a particular mouse that didn't grow at a normal rate. Another technician noticed a mouse that didn't respond normally to the loud noises that occurred when the cages were cleaned—its offspring were found to be susceptible to hereditary deafness. After several experiences like this where unexpected behavior produced important discoveries, the company started holding regular meetings with animal technicians to inquire whether they have spotted anything unusual. These forums for highlighting the importance of unexpected events are called "deviant searches."

CRSS, an architectural firm, developed a special position to repair the problem of theory-based interpretation of evidence. "Most designers love to draw, to make 'thumbnail sketches'," says one manager, but this rush to draw conclusions is often premature. CRSS created a unique job description, the "programmer," to ensure that some members of its design teams were not allowing their own theories to dominate the way they evaluated information from clients. Programmers are not in charge of designing or problem solving, they are in charge of "problem seeking." They are trained to use techniques that help them to resist premature conclusions, and thus listen more carefully to clients. "The experienced, creative [programmer] withholds judgment, resists pre-conceived solutions and the pressure to synthesize...he refuses to make sketches until he knows the client's problem" (Peters, 1992, p. 402).

Often, organizations ensure that individuals weigh information effectively by forcing them to interact with others who might weigh the information differently. One researcher has explored whether training as a scientist cures the problems that other individuals have in evaluating evidence (Dunbar, 1995). The answer is no. For example, scientists, especially young ones, often believe that a single experimental result has just resolved an important problem. However, when Dunbar studied a set of microbiology labs that had been particularly successful, he found that they placed more emphasis on group lab meetings. At these meetings an individual scientist presented his or her results to a variety of skeptical, uninvolved peers. When the individual scientist presented a striking new piece of evidence (e.g., I have detected Enzyme Z in a biological process where it has never been observed before), the individual's peers were typically quite willing to propose alternate ways of interpreting the evidence (e.g., the sample was contaminated with residual Enzyme Z from a prior procedure). In successful labs, even when individual scientists failed to weigh a particular piece of evidence appropriately, their peers did so for them. Moreover, the most successful labs were those that included members with different training and backgrounds. Such "lab meetings" are not limited to successful molecular biology labs; similar meetings take place at venture capital firms where firms decide whether to allocate money to new ventures (Kaplan, 1995).

Individuals Draw Conclusions that are Overconfident and Overly Optimistic

Imagine that individuals have generated a set of hypotheses, collected some new information, and interpreted the relevance of the information for the initial hypotheses. How much confidence should they place in the conclusions they have drawn? If individual learners were adequately cautious, their conclusions would reflect the degree of uncertainty in the data on which they are based. Over the years, research has documented that individuals often express more certainty in their conclusions than is warranted by the facts available to them (or by their actual performance). This kind of problem has been documented extensively in laboratory studies, but also in field studies of individual judgment in a variety of professions, like the civil engineers in the introduction (Griffin & Tversky, 1992; Lichtenstein, Fischhoff, & Phillips, 1982; Russo & Schoemaker, 1992).

Individuals often exhibit a particular kind of overconfidence that we might label a planning fallacy (Buehler, Griffin, & Ross, 1994) or an optimism bias. This optimism bias is pervasive in work environments. Software developers at Microsoft often experience burnout because they "grossly underestimate" how long it will take them to accomplish certain tasks (Cusumano & Selby, 1995, p. 94). Organizations do not always successfully overcome this individual bias. A study of pioneer process plants revealed that the typical plant experienced actual construction costs that were almost double the original projections; similarly, a study of startups showed that more than 80 percent fell short of their projected market share (Davis, 1985). These examples suggest that individuals draw conclusions that underestimate the amount of uncertainty and error in their predictions, but they tend to do it asymmetrically—they rarely *over*estimate a project's cost or time to completion.

Of course, individuals may display an optimism bias because they confront misaligned incentives. Perhaps if engineers correctly estimated the true cost of a new energy plant, decision makers might choose not to build it. However, the real causes of the optimism bias seem to be cognitive, since individuals are overconfident by the same magnitude even in lab experiments that reward accuracy. For example, individuals typically assume that their predictions are more precise than they are. When they are asked to set confidence intervals around a quantity, so that their confidence interval has a 98 percent chance of including the true answer, they are typically surprised by the true answer not 2 percent of the time, but 20 to 50 percent (Lichtenstein, Fischhoff, & Phillips, 1982; Russo & Schoemaker, 1992).

How might organizations repair individual tendencies toward optimism bias and overconfidence? One strategy is to allow individuals to make overconfident predictions, then adjust them overtly. This was the strategy pursued by the engineering profession with its safety factors. Microsoft uses a similar strategy to correct the overly optimistic projections of individual software developers: It has rules about the amount of buffer time that should be added to projects. For reasonably well-understood programming challenges, such as applications programs like Excel and Word, Microsoft typically adds buffer time that constitutes 30 percent of the schedule. However, for operating systems like Windows, where developers must create a system that has to mesh effectively with numerous pieces of hardware and software, Microsoft may add buffer time that reaches 50 percent (Cusumano & Selby, 1995). Similar repairs have evolved in other industries. At one Big Six accounting firm, where teams must prepare formal plans for a consulting engagement, project leaders develop their best estimates of time, expense, and contingency costs, then increase the final number by 15 percent.[12] This repair has evolved despite the fact that this environment provides some incentives to underestimate costs—bids that are too high may not be accepted.

When Microsoft adds buffer time to a schedule, it corrects the predictions of overconfident individuals by overriding them. However, it has also developed procedures that help individual developers decrease their initial level of overconfidence. For example, the company has improved its schedules by requiring developers to create a detailed work plan that specifies which tasks they will perform during specific windows of time. Says one manager, "The classic example is you ask a developer how long it will take him to do something and he'll say a month, because a month equals an infinite amount of time. And you say, 'Okay, a month has 22 working days in it. What are the 22 things you're going to do during those 22 days?' And the guy will say, 'Oh, well, maybe it will take two months.' Even by breaking it down into 22 tasks he realizes, 'Oh, it's a lot harder than I thought'" (Cusumano & Selby, 1995, p. 254).

Some organizations repair overconfidence by forcing individuals to interact with others who are trained to question their conclusions. For example, the Pentagon for many years had what they called the "murder board," a group of experienced officers that reviewed the plans for important missions, with the goal of killing the mission. According to Pentagon lore, the failed Iranian hostage rescue during the Carter years was *not* vetted by this board because high government officials were too concerned about security leaks.[13]

Other organizations have developed norms of frank feedback to ensure that individuals question others' conclusions honestly and openly. In its feature animation unit, Disney regularly holds "Gong Shows" where personnel (including department secretaries) can pitch ideas to a group of senior executives. Gong Shows may attract 40 people who present their idea to the executives and other presenters for three to five minutes. The senior executives are careful to give exceptionally frank feedback at the end of the session, highlighting both good and bad aspects of the presentations. "Somebody may have a great concept, but the story may not be very good. [We can't say] 'Oh, that's fabulous. Great pitch guys!' and when they leave, mumble, 'That was awful!'....We don't pull our punches. [Eventually] people begin to understand that no matter how good, bad, or indifferent the idea, it can be expressed, accepted, and thought about" (McGowan, 1996).

GENERAL DISCUSSION

In this paper we have reviewed the literature on individual learning using a simple framework that considers three broad stages of the learning process. We argued that ideal learners would generate a broad and deep set of hypotheses, test them with a large, unbiased set of information, and draw conclusions in a cautious and balanced way. The psychological literature indicates, however, that real individuals are not ideal learners; they think and act in ways that reduce their ability to learn effectively.

Fortunately, individual learners do not have to go it alone. We have argued that organizations frequently repair the shortcomings of individual learners through the use of sayings, informal routines, and formal procedures. We believe the examples we have offered illustrate the tremendous promise of organizational sources of cognitive repairs.

Nevertheless, we do not think that cognitive repairs will overcome every individual problem. Cognitive repairs are heuristics—like the mental processes they repair, they are pragmatic and often efficient, but also approximate and inexact. For example, they may solve 75 percent of individual shortcomings while incurring only one-third of the costs of optimal procedures (e.g., from economics or statistics). However, they are unlikely to be perfect.

Consider the five whys. It undoubtedly prompts individuals to think more deeply about causes, but it is only a rough heuristic. Why five questions and not three or seven? And which questions? "Problem: He doesn't manage well." (1) Why? He doesn't manage conflict well. (2) Why? He grew up in a dysfunctional family. (3) Why? His parents were alcoholics..." In this example, the answers took an unhelpful detour away from potential solutions sometime around answer 2.

Even when repairs are reasonably effective, they may still leave room for further repair. Consider, for example, the military's partial repair for the fundamental attribution error: "There are no bad troops, only bad officers." This adage may repair tendencies to attribute blame to the people who are closest to a problem (the troops who are on the battlefield); however, it merely focuses attention on another group of people. Thus, it may prevent individuals from fixing systems or procedures that have basic flaws (Cohen & Gooch, 1990). A more effective repair might say, "There are no bad people, only bad systems."

Other repairs may be imperfect because they fix one problem well, but exacerbate others. For example, the Five Cs may help individual loan officers collect more kinds of information but they may create secondary problems. First, by emphasizing character, the Five Cs may provoke the fundamental attribution error. Second, although they expand the set of factors loan officers will consider in a loan decision, they may also institutionalize any tendency that they may have to ignore other potentially relevant factors. Third, they may help loan officers collect information, but they do not necessarily help them interpret it. They seem to indicate that each C should be weighted equally, whereas an ideal statistical model would weigh some Cs more heavily than others.

As these caveats illustrate, cognitive repairs are unlikely to completely repair the shortcomings of individual learners. Thus, when we assess whether a given cognitive repair is successful, we must consider the costs and benefits of the repair. Below, we consider six dimensions that may affect the costs and benefits of repairs, and therefore their success.

Tradeoffs Associated with Successful Repairs

In order to be successful, a cognitive repair must be effective—it must mend some individual shortcoming and improve learning relative to the status quo. To be truly successful, however, a cognitive repair must also be accepted in the organization and actively used. A repair that is not implemented is not a repair.

Cognitive repairs are a kind of innovation, and as such, their use will undoubtedly be affected by many of the characteristics that have previously been mentioned in literatures on diffusion and adoption (Rogers, 1995; Scott, 1995). We will focus on innovation characteristics that are particularly relevant for cognitive repairs. Cognitive shortcomings not only create the need for a repair, they also limit what repairs may succeed.

Below, we consider six dimensions that affect whether a repair will be successful: simple versus complex, domain-specific versus domain-general, familiar versus novel, corrective versus preventative, social versus individual, and top-down versus bottom-up. (We will typically focus on the endpoints of these dimensions, but they should be regarded as continuous rather than dichotomous.) Most dimensions involve tradeoffs. For example, qualities that make a repair more effective in solving an individual shortcoming sometimes reduce the chances that it will be accepted and used by individuals. In the absence of perfectly effective and acceptable repairs, we must recognize and understand the tradeoffs that make repairs more or less successful.

Simple versus Complex

One obvious dimension along which cognitive repairs vary is whether they are relatively simple or complex. Many of the repairs we have discussed in this paper are strikingly simple—they require an individual to remember and apply a procedure that is only a few steps long (e.g., the five whys or the physicians ABCs). In contrast, many of the procedures that are taught as formal repairs in academic environments are quite complex, and involve many stages of sorting, arranging, and calculating (e.g., formal financial or statistical analysis).

Simple repairs have profound advantages over complex repairs. First, they are more likely to be used because the costs are small; individuals will find it easier to learn and implement shorter procedures. By contrast, complex repairs typically require extensive background knowledge and tax basic cognitive resources like attention and memory (Bell, Raiffa, & Tversky, 1988; Nisbett & Ross, 1980). Thus, when individuals encounter a complex repair, they are likely to perceive the costs of learning it as large and immediate, and the benefits of using it as small, uncertain, and delayed.

Second, simple repairs are easier to remember and reconstruct than complex repairs, and this increases the probability that individuals will accurately apply them and accurately transmit them to others. Because complex repairs require individuals to remember a number of stages, they are more likely to be distorted when they are transmitted from individual to individual. This problem will be particularly pronounced in situations that require learning by observation and imitation (DiMaggio & Powell, 1983). Individuals who learn a repair through observation may find it difficult to infer the complete rules of behavior for complex repairs because information about the rules is incomplete, unavailable, or distributed across time in a way that makes learning difficult.

Although simple repairs have profound advantages over complex repairs, they also have some disadvantages. Fundamentally, the tradeoff between simple and complex repairs is a tradeoff between ease of use and accuracy. Complex procedures are often complicated because they attempt to be precise. Simple repairs gain ease of use by sacrificing accuracy. For example, a simple aphorism such as "don't confuse brains and a bull market" suggests the correct direction to adjust one's judgment, but provides no guidance about exactly how much one should discredit the success of an individual trader. To precisely estimate the amount of credit due to brains versus the market, an individual would have to perform a more complex procedure, such as calculating the overall market performance and measuring an individual's performance relative to the dispersion and central tendency of the market.

Domain-Specific versus Domain-General

Cognitive repairs also vary in the range of contexts to which they can be applied (Nisbett, 1992), with some repairs being relatively more domain-specific and some being more domain-general. Domain-specific repairs are tailored narrowly for a specific context (e.g., the Feature Prioritization Process at Motorola or the Five Cs of Credit). Domain-general repairs are described so generally and abstractly that they apply across most judgment tasks (e.g., the Five Whys or most economic or statistical principles).

Domain-specific rules have at least two advantages over domain-general rules. First, individuals find it easier to recognize that a domain specific rule is relevant because the situation itself reminds them to use the rule (e.g., a credit analyst who has learned to think about the Five Cs of Credit will find it difficult to think about lending decisions without considering all five categories of information). Second, individuals may find it easier to apply domain-specific than domain-general rules. Consider, for example a loan officer who is trying to apply a general rule like "calculate the net present value (NPV)" of making the loan. This domain-general rule applies to many more financial decisions than the Five Cs; but it contains no hints about how it should be applied to a loan decision. In contrast, the Five Cs point out specific aspects of the loan decision that might affect the loan's quality. Similarly, securities traders might find it hard to benefit from a domain-general warning against self-serving biases (e.g., "pay attention to situational determinants of success, and don't over-attribute achievement to personal characteristics"). In contrast, they are unlikely to miss the point of a more domain-specific warning not to confuse brains and a bull market.

Although domain-specific rules have advantages, they also have limits. Their specific content will make them more likely to spread within their domain, but it may also prevent them from spreading across domains. For example, engineers have safety factors and software developers have buffer time, but knowing about safety factors does not automatically suggest the need for buffer time. And even a single individual may use a rule effectively in one domain but fail to see its implications for another. Auditors are often quite good at ignoring their preexisting theories about a client's financial health when they investigate auditing problems. However, they are less likely to do so when they confront similar problems outside the auditing domain, even if the problem relates to their other professional activities (Smith & Kida, 1991).

A second limitation of domain-specific repairs is that they are tightly tailored to fit a particular task and environment. Because of this tight fit, they may be less successful than domain-general repairs when the task environment is in flux. A buffer factor designed during a specific period of time—"multiply all time-to-delivery estimates by 1.5"—may lose its effectiveness when technological or economic conditions change. Consider that Microsoft had to develop separate buffer factors to repair overconfidence in applications and operating systems. In general, domain-specific rules will be helpful in companies or divisions where the tasks and environments are stable over time, while domain-general approaches will be helpful in situations where tasks and environments change frequently (e.g., at higher levels in an organization where tasks and decisions are less routine).

A potential method of combining the advantages of domain-specific and domain-general rules may be to give individuals a domain-specific repair and then train them to generalize that repair to other domains (Fong & Nisbett, 1990; Larrick, Morgan, & Nisbett, 1990). Individuals typically find it easier to generalize by analogy from one specific instance to another than to map from general principles down to specifics (Bassok, 1990). For example, people who learn to ignore sunk costs in financial investments correctly recognize that this rule applies to investments of time as well (Larrick, Morgan, & Nisbett, 1990). Similarly, a manager in industry may find it easier to apply the specific military adage about their being "no such thing as bad troops" than to apply a more general lesson about the fundamental attribution error.

Corrective versus Preventative

Cognitive repairs also differ in whether they *prevent* or *correct* the shortcomings of individuals. Corrective repairs intervene during or after a particular cognitive process (e.g., the accounting team that corrects their tendency to overweight vivid exceptions by forcing themselves to consciously classify each exception as major or minor). At the extreme, a corrective repair might only intervene at the very end of a process to correct the overall outcome (e.g., safety factors). Preventative repairs intervene early in a cognitive process before shortcomings have had a chance to act. Microsoft prevents developers from acquiring a biased sample about the speed of their programs by forcing them to develop programs on the same machines used by customers.

Some shortcomings are easier to correct than others. For example, when a shortcoming arises because individuals have the wrong rule, they may not find it difficult to substitute a different rule (Wilson & Brekke, 1994). Trauma physicians may learn to check airway before breathing, and accountants may learn to ignore vivid but minor exceptions. In general, corrective repairs will be appropriate when individuals accept the need for a repair and they understand how to execute the correction.

However, when a shortcoming arises because of some basic cognitive process, organizations may need to intervene more forcefully by bypassing or eliminating the faulty process (Arkes, 1991; Wilson & Brekke, 1994). For example, individuals may find it difficult to generate a broad and independent set of hypotheses because associative memory leads them to consider the same alternatives they previously considered. Theoretically, Motorola could instruct individuals who are developing a new consumer product to "ignore what you've done in the past and approach this problem creatively." However, individuals might find it difficult to ignore their previous solutions. Thus, Motorola prevents the problem by prohibiting them from serving on more than one product development team. Similarly, the Chicago Board of Trade could warn its investigators to discount the answers they receive when they ask leading questions. Instead it prevents individual investigators from asking yes or no questions, and thus ensures that they receive less biased information in the first place.

Familiar versus Novel

Repairs may also vary in the extent to which they are novel rather than familiar. Novel repairs require individuals to change their assumptions or to learn a procedure from scratch. For example the "programmers" at the CRSS architectural firm had to learn to resist their tendency to sketch solutions before evaluating all the information from a client. On the other hand, familiar repairs build on preexisting knowledge (e.g., the CAMEL schema for bank examiners or the ABCs for trauma physicians). These repairs have familiar content-trauma physicians know that they should attend to breathing and circulation, and bank examiners know they should pay attention to capital and earnings. They also have a familiar form—they are organized by a simple acronym that individuals already have as a part of their mental lexicon.

Familiar repairs may be at an advantage over novel repairs because they are less costly to use and their benefits may be more apparent. CAMEL and the ABCs reduce costs by using a familiar acronym to remind individuals to collect types of information that they know they should be collecting. In contrast, the CRSS programmers had to work hard to overcome the behaviors they had learned as architects, and they may have questioned the benefits of the elaborate new procedures they were being taught. Familiar repairs are also less likely to provoke resistance than novel repairs. Anything that requires individuals to throw out old practices or adopt new beliefs may be technically and psychologically difficult.

However, familiar repairs may sometimes be too familiar for their own good. First, they may be less likely to create enthusiasm. If individuals think that a new repair differs only trivially from current practice, they may see no advantage to it. Because individuals often ignore the familiar, would-be change agents often strive to create the perception that their programs are novel and unique (Abrahamson, 1996). Second, familiar repairs may be subject to distortion. If a repair seems partially familiar, individuals may neglect its unfamiliar aspects or force them to mimic the more familiar aspects (a process that psychologists call assimilation). For example, the proper technique for brainstorming requires a specific sequence of steps: first, a creative, idea-generation stage which does not allow criticism, then a stage where ideas are evaluated and selected. Although organizations frequently report that they use brainstorming, careful examination reveals that the organizations are actually engaged in a more familiar activity: a basic business meeting (Zbaracki, in press). The novel aspects of the brainstorming procedure, such as the separation of stages and the "no criticism" rule, are often lost as brainstorming is assimilated to the more familiar practice of the standard business meeting. In the end, only the attractive label remains. In situations where assimilation is a problem, repairs that are novel may be less likely to suffer distortion than repairs that are more familiar because novel repairs do not evoke the preexisting knowledge that leads to assimilation.

Social versus Individual

Many of the cognitive repairs we have considered are social; they work because individuals interact with others (e.g., single-case bore questions in the Deming Prize organization, or the murder board at the Pentagon). Other repairs are individual; individuals apply them to their own learning processes without the intervention of others (e.g., individuals learn to use the Five Whys, and individual investigators at the Board of Trade learn to avoid "yes or no" questions).

In general, we suspect that many successful repairs will be social because individuals may not recognize the need to repair themselves. The very cognitive shortcomings that organizations might hope to repair will make it difficult for individuals to see their own flaws. As we have discussed, individuals tend to retain current assumptions in the face of conflicting data (Klayman, 1995). Also, they interpret events in ways that protect their self-image; they avoid potentially threatening feedback (Frey, 1986; Larrick, 1993) and attribute their poor outcomes to luck or forces outside their control. Although individuals may find it hard to recognize their own biases, they may find it easier to recognize the biases of others. Many of the repairs we document have the feel of (friendly) social gamesmanship. For example, learners at FP&L did not consider the hypothesis that they had been "rucky"—their colleagues considered it for them. Similarly, during weekly microbiology lab meetings, researchers did not have to suggest alternative ways of interpreting their evidence, their peers did so. Social competition among individuals aids the spread of repairs even when individuals are overconfident and believe they would have done just as well without the repair.

Social repairs do have to overcome some disadvantages. For example, individuals may not appreciate others who attempt to repair their biases, and they may dismiss the repair attempts as the product of picky or abrasive personalities. Thus, social repairs may be more successful when an individual understands that his or her tormentors are playing an important, formal role. Individuals may find it easier to entertain an antagonist's critiques when he or she is labeled as a "devil's advocate," or when the individual is appearing before the "murder board." Disney clarified the role of the evaluators and the occasion by establishing its norm of frank feedback and by labeling its tryouts as "The Gong Show."

Eventually, social repairs may be transformed into individual repairs as individuals learn to imitate the patterns of analysis forced on them by others. In order for individuals to learn, they need vivid, immediate feedback. Social encounters are likely to provide a key source of such feedback. For example, when Deming examiners ask single-case-bore questions, or when lab colleagues try to shoot holes in a lab presentation, individual learners may eventually learn to engage in preemptive self-criticism in order to look better during social encounters (Tetlock, 1992). (Many academic papers are better because authors learn to mentally simulate potential reviewer's comments.) Such repairs invoke social forces at two different levels: individuals who anticipate social interaction may be more aware of some of their own shortcomings, and then actual social interaction may overcome additional shortcomings that individuals cannot repair on their own.

Top-Down versus Bottom-Up

Some cognitive repairs originate from "top-down" within an organization. Typically these repairs are deliberately designed and implemented by managers or outside experts. Others arise from bottom-up through informal observation or serendipitous discovery from the people who are doing the work.

The source of the repair is important because it is likely to affect its form. In general, bottom up repairs, such as organizational adages, will be simpler and more domain-specific than top-down repairs designed by technically sophisticated engineers, statisticians, or management gurus. The local origin of bottom-up repairs may also make them feel more familiar and acceptable than top-down repairs. Thus, the origin of a repair will be highly correlated with many of the tradeoffs we have already discussed.

More importantly, the origin of the repair is also likely to affect how potential adopters perceive it. Top-down repairs may be perceived with suspicion or hostility precisely because they originate outside the organization or are imposed from above. Front-line workers may doubt that outsiders like consultants understand their situation well enough to make wise recommendations. When managers suggest a repair, they seem to imply that employees have been performing poorly or cannot be trusted to perform their job correctly. If so, then individuals may resist adopting a repair because of the same kind of self-serving biases we discussed earlier.

Top-down repairs may also be resisted because they will be perceived as driven by politics or fashion and not by the demands of the task. Some top-down repairs may be resisted because they seem too political. Particularly when top-down repairs rely on fixed procedures, they may provoke resistance because individuals may think that they are designed to centralize control or remove individual discretion. Other top-down repairs may be resisted because they seem to be driven by fashion. Institutional theorists contend that organizations adopt new practices for reasons other than greater efficacy. Organizational members may share similar, cynical intuitions (Adams, 1996), and will resist repairs that they see as mere window dressing or as this year's fad. When individuals have trouble recognizing their shortcomings, then they may be particularly likely to attribute top-down repairs to politics or fashion because they will not recognize the repair's true value.

Bottom-up repairs will often benefit from their local, homegrown origins. Local repairs have a meaningful history that makes them memorable and appealing. Even a repair that is potentially threatening, such as "you were rucky," may be more acceptable if organizational members see it as their own invention. Just as lawyers are entitled to tell lawyer jokes, organizational members are entitled to develop self-critical repairs and to convey their insider status by using them. And homegrown repairs evoke a stronger sense of ownership; at the same time that they call attention to a potential shortcoming, they also give the user credit for fixing it.

Conclusions: Successful Repairs

We have considered six different dimensions along which cognitive repairs can be classified. For example, the physician's ABCs are simple, domain-specific, corrective, familiar, individual, and top-down. Although we have suggested some advantages and disadvantages of each endpoint of each dimension, we believe that our discussion suggests at least some preliminary conclusions about successful repairs. For example, because individuals have limited faculties, organizations who wish individuals to learn complex, domain-general repairs will find themselves devoting a great deal of scarce time, money, and effort to ensure that such repairs are learned and used. Similarly, because individuals are overconfident about their own conclusions, they may not spontaneously execute individual repairs to correct their own biases. Based on the advantages and disadvantages we have considered, we suspect that the most successful repairs will be simple, domain-specific, socially administered, and evolved from bottom-up rather than

developed from top-down. We find this conclusion intriguing because it describes repairs that differ sharply from those that are recommended in academic literatures on decision analysis, statistics, and economics.

Implications for Research

Cognitive psychologists often think of people as rugged cognitive individualists, constrained by their own cognitive limitations and a poor environment in which to learn. Cognitive researchers continue to argue over how well individuals actually perform the cognitive tasks they encounter in their lives (see Anderson, 1991, and accompanying commentaries, and the debate between Gigerenzer, 1996; and Kahneman & Tversky, 1996). However, it is important to remember that much of what people do, including much of their cognition, takes place in the context of organizations and other social structures.

Some recent approaches in psychology do explore cultural and social contributions to individual learning. For example, work on "transactive memory" (Wegner, Erber, & Raymond, 1991; Liang, Moreland, & Argote, 1995) shows how individuals reduce their memory limitations by distributing memory-heavy tasks across multiple people. Thus, there can be collective memory that does not reside in any individual. Our concept of organizational repairs is in the same spirit, but it deals with "higher order" cognition: Reasoning and decision making can also be improved through social structure and cultural bootstrapping.

We also believe that organizational psychologists could better understand organizational processes if they thought more about the cognitive processes of the individuals who make up the organization. Research tying individual psychology to organizational behavior certainly has a long and venerable pedigree (March & Simon, 1958), but recently, some researchers have expressed concern that that approach is still underutilized. For example, in the context of institutional diffusion processes, Zucker (1991) has warned that "without a solid cognitive, micro-level foundation, we risk treating institutionalization as a black box at the organizational level, focusing on content at the exclusion of developing a systematic explanatory theory of process" (p. 105). And Miner (1994) warns, "most evolutionary discussions of organizational change discuss routines as though they exist independent of individual human beings" and "evoke images of disembodied entities removed from day-to-day human interaction" (p. 87). Knowledge of individual cognition can be crucial to understanding why things happen as they do in an organization. For example, why do engineering firms use a system by which the best engineers make their best estimates, only to have the firm second-guess them by adding a huge safety factor? Any explanation would be incomplete without an understanding of individual overconfidence. Understanding the abilities and constraints of individuals permits a kind of cognitive archaeology of organizational practices that may allow organizational researchers to better understand why certain rules, norms, and procedures develop in organizations, and why others fail.

Implications for Practice

Managers who think explicitly about cognitive repairs will, we think, be in a better position to foster improvements in their organizations. Managers already think about factors such as incentive systems and information technology as tools to foster learning and innovation. We believe that cognitive repairs will be a useful addition to the toolbox. Managers who consciously consider individual cognition may be able to recognize a larger number of repair opportunities and may target top-down repairs more effectively. Furthermore, they may design more effective repairs if they take a cognitive approach and consider repair dimensions like the six we discussed earlier.

Even when repairs are developed bottom-up rather than top-down, a manager who is informed about individual cognition might have a positive influence. As has been observed by researchers who think about evolutionary approaches to organizations, one of the critical components of organizational learning is to start with a rich and varied pool of alternative practices (Levitt & March, 1988; Miner, 1994). Savvy managers can enhance the pool of alternatives that are available by teaching organization members about the concept of repairs and by encouraging them to identify existing repairs and to seek new repair opportunities. Managers can recognize and reward individuals who discover cognitive repairs, and they can disseminate effective repairs via demonstration, training, communication, and rotation of personnel.

Managers may also find it very important to think about cognitive repairs when they evaluate existing organizational practices. Consider, for example, a new executive who discovers that her development group is split up into six separate teams, each trying to solve the same problem without any communication with each other until well into the process. This might seem like a paradigm of bureaucratic inefficiency. Yet, this is the kind of system that Motorola has found to be effective in generating a broader set of options. Without understanding the potential value of this repair, the new executive might be sadly surprised by the results of streamlining the process.

Managers might also benefit from a cognitive approach because cognitive repairs, like other innovations, may suffer from the law of unintended consequences. A repair that is intended to fix one problem can well end up creating another. Recall, for example, Microsoft's laudable attempt to make the customer more salient to program developers by having them watch live customers attempt to use their products in the usability test lab. The test lab repaired the developer's tendency to be unmoved by cold statistics, but it probably exacerbated their tendency to believe that small samples were reliably representative. In response to seeing one customer in the test lab, developers might waste time altering a feature that would have been okay for most customers. Managers who take a cognitive approach would, we hope, be more likely to avoid unwanted side effects, or at least be in a better position to recognize and cope with them.

Final Words

In contrast to Hamlet's enthusiasm, we argue that there is good reason to be aware of the limitations of individual learners. People are not "infinite in faculties" or "noble in reason." As individuals, we make systematic mistakes that compromise our ability to learn from our experiences and to understand the world. On the other hand, we mortals are not all fools. We are able to form social structures that have the potential to magnify some of our abilities and to minimize some of our failings. In this paper we have concentrated on demonstrating that effective organizational repairs happen. We do not mean to imply that organizations cure any and all individual shortcomings, nor even that organizations always make things better rather than worse. Nevertheless, we do believe that the organizations in which we work can provide us with norms, reminders, protocols, and procedures that help us move beyond our individual shortcomings.

ACKNOWLEDGMENTS

We thank the Graduate School of Business, University of Chicago and the Fuqua School of Business, Duke University, for research support. For helpful comments on this project, we thank participants at Rod Kramer's 1996 conference on Organizational Learning; we also thank Jim March, Cade Massey, Elizabeth McSweeny, Sim Sitkin, Dick Thaler, Elke Weber, and John Wright.

NOTES

1. We primarily want to distinguish problems of incentives from problems of mental process. In this review we will not distinguish between mental errors that arise from "motivated reasoning" and those that arise from "colder" processes.
2. Personal communication, Abhay Joshi.
3. Personal communication, Robert Alan Fisch.
4. Personal communication, Jan Elsbach.
5. Personal communication, Francisco Bayron.
6. Personal communication, Dean Siewert.
7. Personal communication, Justin Scott Bradley.
8. Personal communication, Francisco Bayron.
9. Personal communication, Dean Siewert.
10. Personal communication, Ken Myszkowski.
11. Personal communication, Leo Tucker.
12. Personal communication, Raymond Stukel.
13. Personal communication, John Payne

REFERENCES

Abrahamson, E. (1996). Management fashion. *Academy of Management Review, 21*, 254-285.
Adams, S. (1996). *The Dilbert principle.* New York: Harper Business.

Agostinelli, G, Sherman, S. J., Fazio, R. H., & Hearst, E. S. (1986). Detecting and identifying change: Addition versus deletion. *Journal of Experimental Psychology: Human Perception and Performance, 12*, 445-454.

Anderson, J. R. (1991). Is human cognition adaptive? *Behavioral and Brain Sciences, 14*, 471-517.

Anderson, J.R. (1995). *Cognitive psychology and its implications.* New York: W.H. Freeman & Co.

Arkes, H. R. (1991). Costs and benefits of judgment errors: Implications for debiasing. *Psychological Bulletin, 110*, 486-498.

Baron, J. (1988). *Thinking and deciding.* Cambridge, UK: Cambridge University Press.

Bassok, M. (1990). Transfer of domain-specific problem-solving procedures. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 16*, 522-533.

Bell, D. E., Raiffa, H., & Tversky, A. (1988). *Decision making: Descriptive, normative, and prescriptive interactions.* Cambridge: Cambridge University Press.

Borgida, E., & Nisbett, R. E. (1977). The differential impact of abstract vs. concrete information on decisions. *Journal of Applied Social Psychology, 7*, 258-271.

Buehler, R., Griffin, D., & Ross, M. (1994). Exploring the "planning fallacy": Why people underestimate their task completion times. *Journal of Personality and Social Psychology, 67*, 366-381.

Cheng, P. W., & Holyoak, K. J. (1985). Pragmatic reasoning schemas. *Cognitive Psychology, 17*, 391-416.

Cohen, L. J. (1981). Can human irrationality be experimentally demonstrated? *Behavioral and Brain Sciences, 4*, 317-331.

Cohen, M. D. (1991). Individual learning and oganizational routine: Emerging connections. *Organization Science, 2*, 135-139.

Cohen, E. A., & Gooch, J. (1990). *Military misfortunes: The anatomy of failure in war.* New York: Free Press.

Cusumano, M. A., & Selby, R. W. (1995). *Microsoft secrets.* New York: Free Press.

Davis, D. (1985, March-April). New projects: Beware of false economies. *Harvard Business Review*, 95-101.

Deming, W. E. (1982). *Out of the crisis.* Boston: Massachusetts Institute of Technology Center for Advanced Engineering Study.

DiMaggio, P. J., & Powell W. W. (1983). The iron cage revisited: Institutional isomorphism and collective rationality in organizational fields. *American Sociological Review, 48*, 1,47-160.

Dunbar, K. (1995). How scientists really reason: Scientific reasoning in real-world laboratories. In R. J. Sternberg & J. Davidson (Eds.), *The nature of insight.* Cambridge, MA: MIT Press.

Eisenhardt, K. M. (1989). Agency theory: An assessment and review. *Academy of Management Review, 14* (1), 57-74.

Eppel, D., Argote, L., & Devadas, R. (1991). Organizational learning curves: A method for investigating intra-plant transfer of knowledge acquired through learning by doing. *Organizational Science, 2*, 58-70.

Fischhoff, B., Slovic, P., & Lichtenstein, S. (1978). Fault trees: Sensitivity of estimated failure probabilities to problem representation. *Journal of Experimental Psychology: Human Perception and Performance, 4*, 330-344.

Fiske, S. T., & Taylor, S. E. (1991). *Social cognition* (2nd ed.). New York: McGraw-Hill.

Fletcher, C. (1990). *What cops know.* New York: Pocket Books.

Fong, G. T., & Nisbett, R. E. (1990). Immediate and delayed transfer of training effects in statistical reasoning. *Journal of Experimental Psychology: General, 120*, 34-45.

Forum Corporation (1992). *P.R.O.G.R.E.S.S. tool kit.* Forum Corporation.

Frey, D. (1986). Recent research on selective exposure to information. *Advances in Experimental Social Psychology, 19*, 41-80.

Funder, D. C. (1987). Errors and mistakes: Evaluating the accuracy of social judgment. *Psychological Bulletin, 101*, 75-90.

Gettys, C. F. et al. (1987). An evaluation of human act generation performance. *Organizational Behavior and Human Decision Processes, 39*, 23-51.

Gigerenzer, G. (1996). On narrow norms and vague heuristics: A reply to Kahneman and Tversky (1996). *Psychological Review, 103*, 592-596.

Gilbert, D. T., & Jones, E. E. (1986). Perceiver-induced constraint: Interpretations of self-generated reality. *Journal of Personality and Social Psychology, 50*, 269-280.

Gilbert, D. T., & Malone, P. S. (1995). The correspondence bias. *Psychological Bulletin, 117*, 21-38.

Gnepp, J., & Klayman, J. (1992). Recognition of uncertainty in emotional inferences: Reasoning about emotionally equivocal situations. *Developmental Psychology, 28*, 145-158.

Gregory, W. L., Cialdini, R. B., & Carpenter, K. M. (1982). Self-relevant scenarios as mediators of likelihood estimates and compliance: Does imagining make it so? *Journal of Personality and Social Psychology, 43*, 89-99.

Griffin, D., & Tversky, A. (1992). The weighing of evidence and the determinants of confidence. *Cognitive Psychology, 24*, 411-435.

Heath, C., Knez, M., & Camerer, C. (1993). The strategic management of the entitlement process in the employment relationship. *Strategic Management Journal, 14*, 75-93.

Hilton, D. J. (1995). The social context of reasoning: Conversational inference and rational judgment. *Psychological Bulletin, 118*, 248-271.

Hoch, S. J. (1984). Availability and interference in predictive judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 10*, 649-662.

Hodgins, H. S., & Zuckerman, M. (1993). Beyond selecting information: Biases in spontaneous questions and resultant conclusions. *Journal of Experimental Social Psychology, 29*, 387-407.

Holland, J. H., Holyoak, K. J., Nisbett, R. E., & Thagard, P. T. (1986). *Induction: Processes of inference, learning, and discovery.* Cambridge, MA: Bradford Books/MIT Press.

Hynes, M., & Vanmarcke, E. (1976). *Reliability of embankment performance predictions.* Proceedings of the ASCE Engineering Mechanics Division Specialty Conference. Waterloo, Ontario, Canada: University of Waterloo Press.

Imai, M. (1986). *Kaizen: The key to Japan's competitive success.* New York: McGraw-Hill.

Ishikawa, K. (1982). *Guide to quality control* (2nd revised ed.). Hong Kong: Asia Productivity Organization.

Ishikawa, K. (1985). *What is total quality control? The Japanese way.* Englewood Cliffs, NJ: Prentice-Hall.

Juran, J. M. (1992). *Juran on quality by design: The new steps for planning quality into goods and services.* New York: Free Press.

Kahneman, D., Slovic, P., & Tversky, A. (1982). *Judgment under uncertainty: Heuristics and biases.* Cambridge, England: Cambridge University Press.

Kahneman, D., & Tversky, A. (1996). On the reality of cognitive illusions. *Psychological Review, 103*, 582-591.

Kaplan, J. (1995). *Start-up: A silicon valley adventure.* Boston: Houghton, Mifflin Co.

Keller, P. A., & McGill, A. L. (1994). Differences in the relative influence of product attributes under alternative processing conditions: Attribute importance versus attribute ease of imagability. *Journal of Consumer Psychology, 3*, 29-49.

Koehler, J. J. (1993). The influence of prior beliefs on scientific judgments of evidence quality. *Organizational Behavior and Human Decision Processes, 56*, 28-55.

Klayman, J. (1995). Varieties of confirmation bias. *The Psychology of Learning and Motivation, 32*, 385-418.

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108*, 480-498.

Larrick, R. P. (1993). Motivational factors in decision theories: The role of self-protection. *Psychological Bulletin, 113*, 440-450.

Larrick, R. P., Morgan, J. N., & Nisbett, R. E. (1990). Teaching the use of cost-benefit reasoning in everyday life. *Pyschological Science, 1*, 362-370.

Levitt, B., & March, J. G. (1988). Organizational learning. *Annual Review of Sociology, 14*, 319-40.

Liang, D. W., Moreland, R., & Argote, L. (1995). Group versus individual training and group performance: The mediating factor of transactive memory. *Personality & Social Psychology Bulletin, 21*, 384-393.

Lichtenstein, S., Fischhoff, B., & Phillips, L. D. (1982). Calibration of probabilities: The state of the art to 1980. In A. Tversky, D. Kahneman, & P. Slovic (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 306-334). New York: Cambridge.

Lord, C., Ross, L., & Lepper, M. R. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology, 37*, 2098-2109.

Mahoney, M. J. (1976). *Scientist as subject: The psychological imperative*. Cambridge, MA: Ballinger.

Making quality more than a fad. (1992, May 18). *Fortune*, pp. 12-13.

March, J. G., Sproull, L. S., & Tamuz, M. (1991). Learning from samples of one or fewer. *Organization Science, 2*, 1-13.

March, J. G., & Simon, H. A. (1958). *Organizations*. New York: John Wiley.

McGowan, J. (1996, April 1). How Disney keeps ideas coming. *Fortune*, pp. 131-133.

Milgrom, P., & Roberts, J. (1992). *Economics, organization, and management*. Englewood Cliffs, NJ: Prentice Hall.

Miner, A. S. (1994). Seeking adaptive advantage: Evolutionary theory and managerial action. In J. A. C. Baum & J. V. Singh (Eds.), *Evolutionary dynamics of organizations* (pp. 76-89). New York: Oxford University Press.

Miner, A. S. & Mezias, S. J. (1996). Ugly duckling no more: Pasts and futures of organizational learning research. *Organization Science, 7*, 88-99.

Mullen, B., & Riordan, C. A. (1988). Self-serving attributions for performance in naturalistic settings: A meta-analytic review. *Journal of Applied Social Psychology, 18*, 3-22.

Mynatt, C. R, Doherty, M. E., & Dragan, W. (1993). Information relevance, working memory, and the consideration of alternatives. *Quaterly Journal of Experimental Psychology, 46A*, 759-778.

Newman, J., Wolff, W. T., & Hearst, E. (1980). The feature positive effect in adult human subjects. *Journal of Experimental Psychology: Human Learning and Memory, 6*, 630-650.

Nisbett, R., & Ross, L. (1980). *Human inference: Strategies and shortcomings of social judgment*. Englewood Cliffs, NJ: Prentice-Hall.

Nisbett, R. E. (1992). Reasoning, abstraction, and the prejudices of 20th century psychology. In R. E. Nisbett (Ed), *Rules for reasoning* (pp. 1-12). Hillsdale, NJ: Erlbaum.

Nisbett, R. E., Krantz, D. H., Jepson, D., & Kunda, Z. (1983). The use of statistical heuristics in everyday reasoning. *Psychological Review, 90*, 339-363.

Nonaka, I., & Takeuchi, H. (1995). *The knowledge-creating company: How Japanese companies create the dynamics of innovation*. New York: Oxford University Press.

Odean, T. (1996). Volume, volatility, price and profit when all traders are above average. Working paper. University of California, Berkeley.

Payne, J. W. (1976). Task complexity and contingent processing in decision making: An information search and protocol analysis. *Organizational Behavior and Human Performance, 16*, 366-387.

Perrow, C. (1984). *Normal accidents: Living with high-risk technologies*. New York: Basic.

Peters, T. (1992). *Liberation management: Necessary disorganization for the nanosecond nineties*. New York: Alfred A. Knopf.

Rogers, E. (1995). *The diffusion of innovations* (4th ed.). New York: Free Press.

Rosenthal, E. (1994, May). Invisible injuries. *Discover*, pp. 48-51.

Ross, L., & Nisbett, R. E. (1991). *The person and the situation: Perspectives of social psychology*. New York: McGraw-Hill.

Ross, L. (1977). The intuitive psychologist and his shortcomings. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 10). New York: Academic Press.

Russo, J. E., & Schoemaker, P. J. H. (1992). Managing overconfidence. *Sloan Management Review, 33*, 7-17.

Sapolsky, R. (1997, November). On the role of upholstery in cardiovascular physiology. *Discover*, pp. 58-66.

Scott, W. R. (1995). *Institutions and organizations*. Thousand Oaks, CA: Sage.

Senge, P. M. (1990). *The fifth discipline*. New York: Doubleday Currency.

Slovic, P., Fischhoff, B., & Lichtenstein, S. (1982). Facts versus fears: Understanding perceived risk. In A. Tversky, D. Kahneman, & P. Slovic (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 463-492). New York: Cambridge.

Smith, J. F., & Kida, T. (1991). Heuristics and biases: Expertise and task realism in auditing. *Psychological Bulletin, 109*, 472-489.

Tetlock, P. E. (1992). The impact of accountability on judgment and choice: Toward a social contingency model. *Advances in Experimental Social Psychology, 25*, 331-376.

Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin, 76*, 105-110.

Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology, 5*, 207-232.

Walton, M. (1990). *Deming management at work*. New York: Putnam.

Wegner, D. M., Erber, R., & Raymond, P. (1991). Transactive memory in close relationships. *Journal of Personality and Social Psychology, 61*, 923-929.

What Bill Gates really wants. (1995, January 16). *Fortune*, p. 34.

Wilson, T. D., & Brekke, N. (1994). Mental contamination and mental correction: Unwanted influences on judgments and evaluations. *Psychological Bulletin, 116*, 117-142.

Zbaracki, M. (in press). The rhetoric and reality of total quality management. *Administrative Science Quarterly*.

Zucker, L. G. (1991). Postscript: Microfoundations of institutional thought. In W. W. Powell & P. J. DiMaggio (Eds.), *The new institutionalism in organizational analysis* (pp. 103-107). Chicago: University of Chicago.

Zuckerman, M., Knee, C. R., Hodgins, H. S., & Miyake, K. (1995). Hypothesis confirmation: The joint effect of positive test strategy and acquiescence response set. *Journal of Personality and Social Psychology, 68*, 52-60.