CHAPTER TWELVE

Improving Decision Making

At this point in the book, you may be wondering why human judgment is so terrible. In fact, the situation is not as bad as it seems. Our brains generally serve us well. We are able to perform computational miracles with the three pounds of gray matter between our ears. To pick just two examples, our ability to understand verbal language and to recognize human faces is far beyond that of even the fastest and most powerful computers.

Researchers who study judgment and decision making focus our work on the frailties and shortcomings of human judgment because such examination provides us with the best opportunities to understand the human mind. We learn the most about how we accomplish our goals not by observing successes, but by taking account of failures. When do we confuse one face with another? When do we confuse one word with another? Answers to these questions have helped us understand how our minds process visual and auditory information (Holt & Lotto, 2008; Yovel & Kanwisher, 2005). Just so, the study of judgment biases has revealed a great deal about how people make decisions.

The study of biases is also of immense practical value. Abundant evidence shows that the decisions of smart managers are routinely impaired by biases. Studying how organizations fail can provide useful lessons about what helps them succeed (Bazerman & Watkins, 2004; Perrow, 1984; Ross & Staw, 1986; Sitkin, 1992; Weick, 1993). The good news is that many interventions to improve decision making have emerged in the behavioral decision research literature, and many of these interventions have been developed and have succeeded in the real world.

One story of an effective decision-changing process appears in Michael Lewis's 2003 book *Moneyball* and the 2011 film adaptation. Lewis tells the story of how Billy Beane, the general manager of the Oakland Athletics, transformed his baseball team by questioning the intuition of baseball professionals. From 1999, when Beane took over as a general manager of the Oakland Athletics, through 2002, the team achieved a truly amazing record. The year Beane took over, the team ranked eleventh of fourteen in the American League in terms of payroll yet placed fifth out of fourteen in wins. In both the 2000 and 2001 seasons, the Athletics ranked twelfth in payroll and second in wins in the American League. In 2002, they were twelfth in payroll and first in wins in the league. Over this four-year period, the team had the second-best record in Major League Baseball with

one of the two smallest payrolls in the entire league. The players earned less than a third of the amount earned by the New York Yankees yet won more games than the Yankees.

How did the Athletics achieve this success? The simple answer is that manager Billy Beane, with the help of Paul DePodesta, a recent Harvard economics graduate, realized that the intuition of baseball executives was limited and systematically biased, and that their perceived "wisdom" nonetheless had been incorporated into personnel management in ways that created enormous inefficiencies. Lewis (2003) argues that baseball executives were consistently guilty of three mistakes. First, they overgeneralized from their personal experiences. Second, they were overly influenced by players' recent performances. Third, they were overly influenced by what they saw with their own eyes, when players' multiyear records provided far better data.

More broadly, Beane and DePodesta found that expert intuition in baseball systematically overweighted some variables and underweighted other variables. The results made it clear that, in baseball, statistics have outperformed the experts. After allowing intuition to rule decision making in baseball for over one hundred years, teams are only now replacing their "experts" with nerds who know how to run regression equations. In Lewis's (2003) words, "the market for baseball players was so inefficient, and the general grasp of sound baseball strategy so weak, that superior management could run circles around taller piles of cash." Following Beane's success, many teams tried to hire DePodesta as their general manager, and most teams learned to rely more heavily on statistical analysis to predict players' future performance (Schwarz, 2005).

The story of the Athletics' success raises some interesting questions. Why did it take so long for rationality to enter into decision making in baseball? To what extent are managers in other industries still relying on false expertise when better strategies exist? As Thaler and Sunstein (2003) note in their insightful review of *Moneyball,* baseball professionals are not stupid, but they are human. Like all of us, they have tended to rely on simple heuristics, traditions, and habits, which in turn created the conventional wisdom that governed baseball for over a century. It takes time, effort, and courage for an organization to move from relying on faulty intuition to carefully assessing data and using appropriate statistical techniques.

Lewis (2003) argues that the mistakes documented in Major League Baseball are probably more severe in other industries. After all, the sport of baseball is full of excellent, reliable data. Thaler and Sunstein (2003) compare the tendency of baseball executives to overlook a wealth of statistics to the tendency of personnel managers to base hiring decisions on their "gut" reactions to job interviews rather than on the hard data available on applicants. Executives tend to trust their intuitive reactions to interviews, despite extensive research showing that interviews provide little predictability about future performance. Thaler and Sunstein (2003) argue for personnel selection based on real performance predictors (grades, test scores, past company performance, etc.) rather than on intuition gathered from interviews.

In this chapter, we argue that most organizations have the opportunity to significantly increase the effectiveness of their decision-making processes. We will not argue that executives are lacking in intelligence. Rather, like baseball executives, most professionals make decisions that fall short of objectively rational behavior and do so in specific and systematic ways. The critical question is: What can we do to correct these deficiencies? This concluding chapter examines seven concrete and complementary strategies for making better decisions: (1) use decision-analysis tools, (2) acquire expertise, (3) debias your judgment, (4) reason analogically, (5) take an outsider's view, (6) understand biases in others, and (7) nudge people toward wiser and more ethical decisions.

STRATEGY 1: USE DECISION-ANALYSIS TOOLS

Because we do not make optimal decisions intuitively or automatically, when decision quality really matters, it makes sense to rely on procedures that can help direct us toward more optimal decisions. The field of study that specializes in giving this sort of prescriptive decision advice is generally called *decision analysis*. A number of books have distilled the field's wisdom to provide useful guides for making decisions (for example, see Goodwin, 1999; Hammond, Keeney, & Raiffa, 1999). These approaches usually require you to quantify both your preferences and the value you place on each of the various decision options. Rational decision-making strategies also require you to be specific about the probabilities associated with uncertain future outcomes.

Decision analysis usually guides decision making using the logic of *expected value*. To compute an option's expected value, you multiply its value by its probability. So, for instance, to compute the dollar value of a lottery ticket, you would need to multiply the dollar value of its payout by the probability of receiving that payout. Because the expected value of lottery tickets is almost always less than it costs to buy them, purchasing lottery tickets is usually not a good use of your money. When a decision has multiple dimensions, such as a choice between two houses, one that is expensive and newly renovated and another whose price is more reasonable but that requires more work, the decision usually requires some sort of multi-attribute utility computation. This computation forces the decision maker to weigh her willingness to spend money against her willingness to perform home improvement work.

Often, however, businesses need to make a series of similar decisions over and over. For instance, corporations need to decide which applicants to hire. Executives need to decide which employees to promote and how big each employee's bonus should be. Bank loan officers need to decide whether to extend credit to loan applicants. Venture capitalists need to decide whether to fund an entrepreneur's new venture. These complex decisions can be guided by the use of a linear model.

What is a Linear Model?

A linear model is a formula that weights and adds up the relevant predictor variables in order to make a quantitative prediction. As an example, when his older

son was five, Don asked the boy's pediatrician to predict how tall Josh would grow to be. The pediatrician offered a simple linear model in response. She said that a child's adult height is best predicted with the following computation. First, average the parents' heights. Second, if the child is a boy, add two inches to the parents' average. If the child is a girl, subtract two inches from the parents' average. Innumerable linear models such as this exist to help us make informed predictions. A linear model called PECOTA, for instance, helps baseball teams predict players' future performances using data such as their prior performances, ages, heights, and weights (Schwarz, 2005). There is even a company that uses a secretive linear model to help movie studios predict how much money their movies will earn (Gladwell, 2006).

Why Linear Models Can Lead to Superior Decisions

Researchers have found that linear models produce superior predictions than experts across an impressive array of domains. In addition, research has found that more complex models produce only marginal improvements above a simple linear framework. Dawes (1979) argues that linear models are superior because people are much better at selecting and coding information (such as what variables to put in the model) than they are at integrating the information (using the data to make a prediction). Einhorn (1972) illustrates this point in a study of physicians who coded biopsies of patients with Hodgkin's disease and then made an overall rating of disease severity. The individual ratings were not able to predict the survival time of the patients, all of whom died of the disease. However, the variables that the physicians selected to code did predict survival time when optimal weights were determined with a multiple regression model. The doctors knew what information to consider, but they did not know how to integrate this information consistently into valid predictions.

In addition to having difficulty integrating information, we are also inconsistent. Given the same data, we will not always make the same decision. Our judgment is affected by mood, subjective interpretations, environment, deadlines, random fluctuations, and many others nonstable characteristics. In contrast, a linear model will always make the same decisions with the same inputs. Thus, such models capture the underlying policy that an expert uses while avoiding the expert's random error. Furthermore, experts are likely to be affected by certain biases triggered by specific cases. In contrast, linear models include only the actual data that are empirically known to have predictive power, not the salience or representativeness of that or any other available data. In short, linear models can be programmed to sidestep biases that are known to impair human judgment.

Such bias is common in financial decisions, corporate personnel decisions, bank loan decisions, and routine purchasing decisions. In each of these domains, the decision maker must make multiple routine decisions based on the same set of variables—a task well suited to a linear model. Such models allow the organization to identify the factors that are important in the decisions of its experts. Thus,

independent of their superior predictive powers, the feedback and training opportunities provided by linear models make them a valuable managerial tool.

Why We Resist Linear Models

While evidence amply supports the power of linear models, such models have not been widely used. Why not? Resistance to them is strong. Some have raised ethical concerns, such as this one described by Dawes:

> I overheard a young woman complain that it was "horribly unfair" that she had been rejected by the Psychology Department at the University of California, Santa Barbara, on the basis of mere numbers, without even an interview. "How could they possibly tell what I'm like?" The answer is they can't. Nor could they with an interview.

Dawes argues that decision makers demonstrate irresponsible conceit in believing that a half-hour interview leads to better predictions than the information contained in a transcript covering three-and-a-half years of work and the carefully devised aptitude assessment of graduate board exams.

Now consider the response that Max received when he asked a well-known arbitrator to make a number of decisions as part of a study of arbitrator decision-making processes:

> You are on an illusory quest! Other arbitrators may respond to your questionnaire; but in the end you will have nothing but trumpery and a collation of responses which will leave you still asking how arbitrators decide cases. Telling you how I would decide in the scenarios provided would really tell you nothing of any value in respect of what moves arbitrators to decide as they do. As well ask a youth why he is infatuated with that particular girl when her sterling virtues are not that apparent. As well ask my grandmother how and why she picked a particular "mushmelon" from a stall of "mushmelons." Judgment, taste, experience, and a lot of other things too numerous to mention are factors in the decisions (Bazerman, 1985).

In contrast with this arbitrator's denial of the possibility of systematically studying decision processes, research in this area does show that linear models are capable of capturing his decision-making model (or his grandmother's choice of mushmelon).

Another argument commonly made against decision-analysis tools such as linear models is that they rule out the inclusion of intuitions or gut feelings. In an apocryphal story, Howard Raiffa was on the faculty at Columbia and received an offer from Harvard. According to the story, he visited his dean at Columbia, who was also his friend, and asked for help with his decision. Sarcastically, the dean, borrowing from Raiffa's writings on decision analysis, told Raiffa to identify the relevant criteria, weight each criterion, rate each school on each criterion, do the arithmetic, see which school had the best overall score, and go there. Supposedly,

Raiffa protested, "No, this is a serious decision!" While he enjoys this story, Raiffa says it simply isn't true. The more important the decision is, he continues to believe, the more important it is to think systematically about it.

Finally, people sometimes argue that the use of linear models will require difficult changes within organizations. What will bank loan officers or college admissions officers do when computers make the decisions? Such concerns express the fear that people are not necessary for linear models to make decisions. In fact, people play a crucial role in models. People decide which variables to put into the model and how to weight them. People also monitor the model's performance and determine when it needs to be updated. Nevertheless, resistance to change is natural, and resistance to the use of linear models is clearly no exception. Overcoming a bias against expert-based, computer-formulated judgments is yet another step you can take toward improving your decision-making abilities. We will now look more closely at two domains in which evidence shows that linear models can lead to better organizational outcomes: graduate-school admissions decisions and hiring decisions.

Improving Admissions Decisions

The value of using linear models in hiring, admissions, and selection decisions is highlighted by research on the interpretation of grades (Moore, Swift, Sharek, & Gino, 2010). There are substantial differences in the grading practices of colleges, even between institutions of similar quality and selectivity. It turns out that students from colleges with more lenient grading are more likely to get in to graduate school, even after controlling for the quality of the institution and the quality of its students. In one study, due to a variant of the representativeness heuristic called the *correspondence bias* (Gilbert & Malone, 1995), graduate schools mistook the high GPAs of alumni from lenient-grading institutions as evidence of high performance. The correspondence bias describes the tendency to take others at face value by assuming that their behavior (or their GPAs) corresponds to their innate traits. The researchers found that this bias persisted even when those making the admissions decisions had full information about different institutions' grading practices. It seems that people have trouble sufficiently discounting high grades that are due to lenient grading.

By contrast, it would be easy to set up a linear program to avoid this error. Indeed, Dawes (1971) did just that in his work on graduate-school admissions decisions. Dawes used a common method for developing his linear model: he first modeled the admission decisions of a four-person committee. In other words, he systematically analyzed how the committee made its admissions decisions, relying on three factors: (1) Graduate Record Examination scores, (2) undergraduate GPA, and (3) the quality of the undergraduate school. Dawes then used the variable weightings he obtained from modeling the experts in a linear model to predict the average rating of 384 other applicants. He found that the model could be used to rule out 55 percent of the applicant pool without ever rejecting an applicant that the selection committee had in fact accepted. In addition, the linear

model was better than the committee itself in predicting future ratings of the accepted and matriculated applicants by faculty! In 1971, Dawes estimated that the use of a linear model as a screening device by the nation's graduate schools could result in an annual savings of about $18 million in professional time. Adjusted for today's dollars and the current number of graduate-school applications, that number would easily exceed $500 million. And this figure neglects many larger domains, including undergraduate admissions and corporate recruiting.

Improving Hiring Decisions

Hiring decisions are among the most important decisions an organization can make. Virtually every corporation in the world relies on unstructured, face-to-face employment interviews as the most important tool for selecting employees who have passed through an initial screening process. The effectiveness of employment interviews for predicting future job performance has been the subject of extensive study by industrial psychologists. This research shows that job interviews do not work well. Specifically, employment interviews predict only about 14 percent of the variability in employee performance (Schmidt & Hunter, 1998). In part, this figure is so low because predicting job performance is difficult and few tools do it well. Yet some assessment tools do predict performance substantially better than the unstructured interview, and at a substantially lower cost.

So why do people continue to believe so strongly in employment interviews? Managers' robust faith in the value of interviews is the result of a "perfect storm" of cognitive biases:

- *Availability*: Interviewers may think they know what constitutes superior employee performance, but their information is highly imperfect. Few companies bother to collect useful data on the attributes that employees need to succeed within specific positions or within the broader organization. As a result, managers must rely on their intuitions to determine whether or not a job candidate has the qualities needed for success.
- *Affect heuristic*: People make very quick evaluations of whether they like others or not based on superficial features, such as physical attractiveness, mannerisms, or similarity to oneself (Ambady, Krabbenoft, & Hogan, 2006; Ambady & Rosenthal, 1993). Managers rarely revise these first impressions in the course of an employment interview (Dougherty, Turban, & Callender, 1994). Managers sometimes claim that interviews allow them to assess a potential candidate's "fit" with the firm, but this assessment is usually not based on systematic measurement of a candidate's qualities and is little more than the interviewer's intuitive, affective response.
- *Representativeness*: Intuition also leads managers to believe that if a person can speak coherently about her goals, the organization, or the job, then she will perform well at the job. For most jobs, however, interview performance is weakly related to actual job performance. Extroverted, sociable, tall, attractive, and ingratiating people often make more positive interview impressions than

others. However, these traits are often less critical to job performance than other, less immediately observable traits, such as conscientiousness and intelligence.

- *Confirmation heuristic*: After interviewing a number of people for a position and hiring one of them, managers only learn about the performance of the person selected. Without knowing whether that person is performing better than the rejected applicants would have, managers lack the data they would need to assess whether their selection mechanisms are effective (Einhorn & Hogarth, 1978).

What is a better alternative to face-to-face, unstructured employment interviews? A number of other selection tools are available, most of which are less expensive to implement than interviews, including simple intelligence tests. But if organizations insist on conducting interviews, they ought to use structured ones in which all job candidates are reviewed by the same set of interviewers and in which each interviewer asks the same questions of each candidate (Schmidt & Hunter, 1998). In addition, interviewers' quantitative assessments ought to be just one component fed into a linear model, along with intelligence measures, years of relevant work experience, and so on.

STRATEGY 2: ACQUIRE EXPERTISE

Many of the biases we have examined in this book were identified in experiments with student participants who were not rewarded for accurate performance and who were making decisions in task domains unfamiliar to them. Thus, one optimistic possibility is that experts or experienced decision makers facing important real-world decisions might be far less affected by biases than most research participants. Does this book unfairly exaggerate the prevalence of judgment biases? This is certainly an important question, since experience and expertise might be useful tools for improving decision making.

Some researchers believe that the process of improving judgment will occur naturally as individuals receive feedback about their past decisions. This view is represented by Kagel and Levin (1986, p. 917) in their analysis of the winner's curse in competitive bidding discussed in Chapter 4:

> Given sufficient experience and feedback regarding the outcomes of their decisions, we have no doubt that our experimental participants, as well as most bidders in "real world" settings, would eventually learn to avoid the winner's curse in any particular set of circumstances. The winner's curse is a disequilibrium phenomenon that will correct itself given sufficient time and the right kind of information feedback.

In fact, Kagel and Levin (1986) do show a reduction in the winner's curse in the auction context as the market (but not necessarily specific players) "learns" over time. However, much of this learning can be attributed to the phenomenon in

which the most aggressive bidders go broke and drop out of the market. Additional learning occurs by observing the consistent losses being suffered by "winners" in the auction.

Clearly, life experiences help us to improve numerous skills and abandon many bad habits. Unfortunately, our judgmental distortions might not be among them. Tversky and Kahneman (1986) have argued that basic judgmental biases are unlikely to correct themselves over time. Responsive learning requires accurate and immediate feedback, which is rarely available in the real world because:

> (i) outcomes are commonly delayed and not easily attributable to a particular action; (ii) variability in the environment degrades the reliability of feedback. . . ; (iii) there is often no information about what the outcome would have been if another decision had been taken; and (iv) most important decisions are unique and therefore provide little opportunity for learning (see Einhorn and Hogarth, 1978) . . . any claim that a particular error will be eliminated by experience must be supported by demonstrating that the conditions for effective learning are satisfied (pp. s274–s275).

Even if accurate and immediate feedback is available in a given situation, we face another crucial challenge: we are likely to misremember our own forecasts (Meyvis, Ratner, & Levav, 2010). We often anchor to current states and fail to accurately recall our prior predictions. Thus, it is common for us to underestimate the extent to which our prior predictions deviated from actual outcomes, and this underestimation leads to us inadequately learn from prior experience (Morris & Moore, 2000).

Using the "Acquiring a Company" problem described in Chapter 4, Ball, Bazerman, and Carroll (1991) tested the ability of individuals to learn to avoid the winner's curse by incorporating the decisions of others into their decision making. Participants in this experiment played for real money, played in 20 trials, and were given full feedback immediately after each trial based on a random determination of the value of the firm up for sale; in addition, they could observe changes in their asset balance (which virtually always went down). Thus, when compared to the limitations cited by Tversky and Kahneman, ideal conditions existed for learning from past mistakes. The only limitation that was not eliminated—namely, the variability of the environment (ii above)—is a natural part of the winner's curse phenomenon. Thus, we were able to look at whether or not the ability to consider the cognitions of the other party in a bilateral negotiation problem can be learned in a highly favorable environment.

Remembering that $0 is the correct answer and that $50 to $75 is the answer typically obtained when decision makers ignore the cognitions of others, examine the mean bids across the 20 trials in Figure 12.1. Across the 20 trials, there is no obvious trend indicating that participants learned the correct response. In fact, only five of 72 participants from a leading MBA program learned over the course of the trials. Our general conclusion? Individuals are unlikely to overcome the winner's curse simply through experience or feedback.

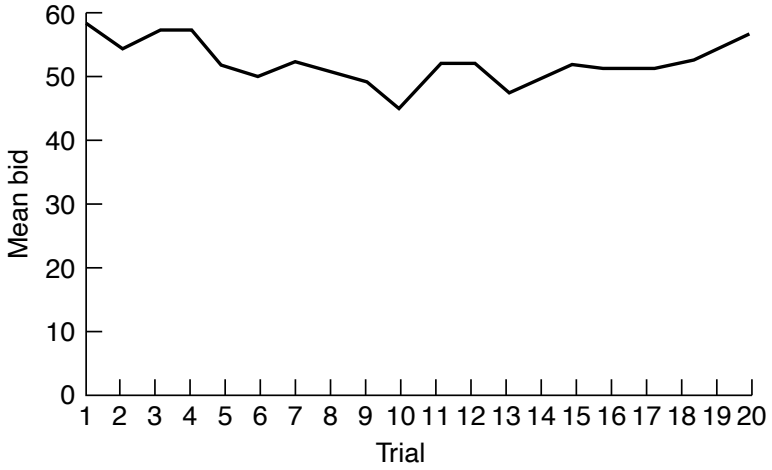

Figure 12.1 Mean Offers Across 20 Trials of the "Acquiring a Company" Problem.

This evidence paints a pessimistic picture of the idea that experience will cure the decision biases identified in this book. In fact, Bereby-Meyer and Grosskopf (2008) documented that even hundreds of trials do not lead most study participants to solve the Acquiring a Company problem. This evidence is consistent with the documentation of extensive bias in decision making by actual investors, real-estate agents, medical doctors, and numerous other "expert" groups. Neale and Northcraft (1989) proposed that biased decision-making outcomes could be eliminated or ameliorated through the development of expertise. While we often think of experience and expertise as closely related, Neale and Northcraft defined experience simply as repeated feedback. By contrast, they assert that expertise results when individuals develop a "strategic conceptualization" of what constitutes a rational decision-making process and learn to recognize the biases that limit rationality.

Neale and Northcraft's experience/expertise distinction is highly relevant to the question of whether or not experienced decision makers can benefit from the study of decision making. Northcraft and Neale's (1987) study of anchoring and adjustment among real-estate agents suggests that experienced decision makers can be very biased. In addition, while most "effective decision makers" are successful in a specific domain, experience without expertise can be quite dangerous when it is transferred to a different context or when the environment changes. Evidence from Chapter 2 suggests that as the amount of ignorance increases, individuals become more overconfident regarding their fallible judgment.

If you think that experience should help negotiators do a better job of understanding the other side's reservation price, think again. Larrick and Wu (2007) find that, when it comes to estimating the size of the bargaining zone, experience will only help us correct one type of error: overestimation of the bargaining zone's size. When you think the bargaining zone is much bigger than it is, your negotiating counterpart will help you identify and correct your error by

refusing to agree to deal at the price you propose. When, on the other hand, you underestimate the size of the bargaining zone, you will end up offering the other side more than was necessary. Though she probably will be anxious to accept your offer, she may try to get you to concede a bit more first, so that you will think that your offer is close to her reservation price. This type of experience will generally lead negotiators to believe that bargaining zones are smaller than they actually are and that they need to make more generous offers to their negotiating opponents.

Stressing the drawbacks of relying on experience for knowledge, Dawes (1988) notes that Benjamin Franklin's famous quote "experience is a dear teacher" is often misinterpreted to mean "experience is the best teacher," when in fact Franklin was using "dear" as a synonym for expensive. After all, the quote continues, "yet fools will learn in no other [school]." Dawes writes,

> Learning from an experience of failure . . . is indeed "dear," and it can even be fatal. . . . moreover, experiences of success may have negative as well as positive results when people mindlessly learn from them. . . . People who are extraordinarily successful—or lucky—in general may conclude from their "experience" that they are invulnerable and consequently court disaster by failing to monitor their behavior and its implications.

Or in the words of Confucius: "By three methods we may learn wisdom: First, by reflection, which is noblest; Second, by imitation, which is easiest; and third, by experience, which is the bitterest."

This view of experience reiterates the comparative value of gaining a conceptual understanding of how to make a rational decision rather than simply depending upon the relatively mindless, passive learning obtained via experience. Expertise requires much more than the unclear feedback of uncertain, uncontrollable, and often delayed results. Rather, it necessitates constant monitoring and awareness of our decision-making processes. The final benefit of developing a strategic conceptualization of decision making concerns transferability. If you ask experienced decision makers for the secrets of their success, they routinely insist that their skills have developed over years of observation and experience that cannot be taught. This obviously reduces their ability to pass on their knowledge to others. Thus, experience without expertise limits the ability to transfer knowledge to future generations.

A key element of developing a strategic conceptualization of decision making is to become aware of the many biases in individual and group contexts that we have discussed in Chapters 1 through 11. However, awareness is just one step in the process. Another strategy, debiasing, is the topic of the next section.

STRATEGY 3: DEBIAS YOUR JUDGMENT

Debiasing refers to a procedure for reducing or eliminating biases from the cognitive strategies of the decision maker. Fischhoff (1982) proposed four steps that decision-making teachers or trainers can follow to encourage their students to

make wiser judgments: (1) offer warnings about the possibility of bias, (2) describe the direction of the bias, (3) provide a dose of feedback, and (4) offer an extended program of training with feedback, coaching, and whatever else it takes to improve judgment. Fischhoff also argues that debiasing is an extremely difficult process that must be closely monitored and guided by a psychological framework for change. For example, research on the hindsight bias (Fischhoff, 1977), described in Chapter 3, has shown that even when the bias is explicitly described to participants and they are instructed to avoid it, the bias remains.

In contrast, a review by Larrick (2004) paints a rosier picture of our ability to overcome bias through training. Yet Larrick (2004) also notes that most successful debiasing strategies tend to be context- and bias-specific; training and testing must be closely linked and must occur in close time proximity. For example, research on the overconfidence bias has found that intensive, personalized feedback is moderately effective in improving judgment (Lichtenstein & Fischhoff, 1980), but only in the short term. Occasionally, a broader effect of training has been documented. For example, simply encouraging people to "consider the opposite" of whatever they are deciding reduces overconfidence, hindsight, and anchoring effects (Larrick, 2004; Mussweiler, Strack, & Pfeiffer, 2000). Larrick (2004) also highlights the partial debiasing success of using groups instead of individuals, training in statistical reasoning, and making people accountable for their decisions (Lerner & Tetlock, 1999).

Based on Lewin's framework outlined in Chapter 1, Fischhoff's debiasing research, Larrick's review, and our own judgment-training programs with MBA and executive students, this section makes specific suggestions for debiasing judgment.

Unfreezing

Chapter 1 noted that many behaviors at the individual, group, and organizational levels are ingrained, or part of a standard repertoire, and are therefore quite difficult to change. Factors that inhibit individuals from changing their behavior include satisfaction with the status quo, risk aversion, and a preference for the certain outcomes of known behavior to the uncertain outcomes of innovative behavior. For improved decision making to occur and continue over time, an explicit "unfreezing" process of ingrained thinking and behaviors must take place. For at least three key reasons, unfreezing old strategies is crucial to changing the decision-making processes of individuals.

First, individuals will have typically relied on their current intuitive strategy for many years. To want to change would be to admit that past strategies were flawed, and this realization is likely to be disturbing. Thus, individuals may be motivated to avoid the disconcerting truth about their judgmental deficiencies.

Second, individuals who have achieved a certain level of professional success (such as students in MBA and executive education programs) are likely to have received positive reinforcement for many of their past decisions. According to the basics of reinforcement theory, individuals tend to continue behaviors that are positively rewarded. For example, because many successful executives rise to the

top using intuitive strategies, they tend to resist information indicating that their judgment is systematically deficient in some demonstrable manner.

A third, related point has to do with balance theory (Heider, 1958), which suggests that individuals try to manage their cognitions into a consistent order. For successful managers, the notion that "there is something fundamentally wrong with my decision-making processes" clashes with their awareness of their success. The belief "I am currently an excellent decision maker" is much more harmonious with the notion of success; therefore, according to balance theory, that cognition is more likely to dominate.

Overall, a pattern emerges of an intelligent manager who has multiple reasons for believing in the high quality of his or her decision-making processes and resisting any change to his or her intuitive strategies. Most successful people will be motivated to view their intuition as a talent rather than a handicap. In fact, this book has provided substantial evidence that there is significant room for improvement in the intuitive strategies of even the brightest, most successful managers. Thus, we conclude that improving on intuition is an important activity for successful managers to attempt, but that cognitive resistance to change is a predictable pattern.

This book has sought to create changes in your judgment by exposing you to concrete evidence that leads you to question your current strategies. The quiz-and-feedback format was designed specifically to unfreeze your decision-making processes. Most readers make a substantial number of mistakes on these items and are then ready to learn where they went wrong and how they could have performed better. This format unfreezes the notion that your decision-making processes do not require improvement. As you begin to question your current strategies, you become receptive to alternatives. In other cases (such as the dollar auction), vivid examples were intended to unfreeze your thinking by leading you to identify with individuals who fell victim to judgmental deficiencies.

Change

Once you has unfrozen past behaviors, you will become willing to consider alternatives. The next stage consists of making the change itself. However, change is far from guaranteed; internal resistance is likely, causing you to continually reassess the desirability of change. There are three critical steps to changing your decision-making process: (1) clarification of the existence of specific judgmental deficiencies, (2) explanation of the roots of these deficiencies, and (3) reassurance that these deficiencies should not be taken as a threat to your self-esteem.

The first step consists of abstracting from the concrete example that was used for unfreezing to identify the more general bias that exists. In addition, for the bias to have face validity to you, an explanation of why the bias exists is necessary; this often consists of clarifying the heuristic or phenomenon that underlies the bias. Finally, this information may be threatening enough to increase the resistance that you partially overcame in the unfreezing stage. Thus, it is critical that you understand that virtually everyone is subject to judgment biases and that having them does not imply that you are a poor decision maker, but simply human.

Perhaps the most general-purpose debiasing strategy is what Lord, Lepper, and Preston (1984) call "consider the opposite." They advise us to play devil's advocate with ourselves and to think about reasons our tentative conclusions could be wrong. This strategy is obviously most useful for counteracting the confirmation trap—the tendency to seek out information that supports our chosen point of view while overlooking disconfirming evidence. Baron (1994) has given more specific advice. He suggests that, when assessing any piece of data, you should do two things. First, ask yourself: "How likely is a yes answer, if I assume that my hypothesis is false?" For instance, imagine you are considering investing money in a friend's new business idea. You take it as a good sign that his business plan projects he will turn a profit in one year. The hypothesis you've been entertaining is that this is a good investment for your money. But what if you assume the hypothesis is false: that this investment is a terrible idea because it puts both your money and your friendship in peril? Is it possible that your friend came up with a plausible business plan but that his chances of success are not particularly great?

Second, try to think of alternative hypotheses, then choose a test most likely to distinguish them. Could you devise a test that could tell whether your friend's plan was actually a viable one? Maybe the fact that he has had trouble getting start-up funding from banks or venture capitalists is a sign that his business plan doesn't stack up that well against those of other aspiring entrepreneurs. This process is useful not only for counteracting the confirmation bias but also for reducing overconfidence. Admittedly, it's not always fun to consider ways in which we might be wrong, but this is a crucial step when sound decisions and accurate judgments are more important than ego gratification.

Refreezing

After we make a positive change, it is tempting to revert to past practices and bad habits. The old biases still exist and can be easily and even accidentally used. Meanwhile, the new procedures are foreign and must develop into intuitive strategies, a process that takes place with practice over time. As you consciously use new strategies in multiple applications, these strategies slowly become second nature, taking the place of old patterns. However, frequent application and overviews of past training are necessary if change is to last.

For refreezing to occur, you must continue to examine your decisions for bias long after you have finished this book. You should schedule routine "checkups" to evaluate your recent important decisions—those you made on your own, as a negotiator, and as a member of a group—while remaining aware of the limits of your judgment.

STRATEGY 4: REASON ANALOGICALLY

Analogical reasoning, or the process of abstracting common lessons from two or more situations, turns out to be a remarkably simple debiasing approach (D. Gentner, G. Loewenstein, & L. Thompson, 2003a; Loewenstein, Thompson,

& Gentner, 1999; L. Thompson, D. Gentner, & J. Loewenstein, 2000). Research shows that people learn far more from cases, simulations, and real-world experiences when they are able to take away an abstract form of the learning message. In the context of learning to negotiate through simulations, Gentner, Loewenstein, & Thompson (2003b) found that greater debiasing occurred among participants when they took part in two exercises that had the same lesson and were asked how the two simulations were related than when they assessed the same two exercises and were asked to explain the lesson of each one. When people learn from one episode at a time, they too often focus on superficial characteristics of the situation and assume that the message applies only to the specific context of the decision (such as learning how to buy a house). By contrast, the process of abstracting similar lessons from two episodes (such as learning to overcome the mythical fixed pie of negotiation following a house purchase and a workplace negotiation) creates more generalizable insight.

By assessing participants' performance on a third task, Gentner, Loewenstein, and Thompson (2003a) have demonstrated evidence of debiasing decision-making and negotiation behavior through this type of analogical reasoning. They have replicated their research conclusions across a number of studies, many involving executives and consultants. Thompson, Gentner, & Loewenstein (2000) claim that when people make a comparison, they focus on the similarities between examples, whose common structure becomes more transparent. Identifying the common structure—the principle shared by both examples—helps the learner form a schema that is less sensitive to the irrelevant surface or context features of the particular examples. Such an abstract principle is more likely to be transferred to new situations with different contexts than a principle that is not abstracted from its original context. These impressive findings on the effectiveness of analogical reasoning open up important new directions for debiasing research and offer guidance on how to use cases and simulations to maximize generalizable learning.

Building on Thompson et al.'s analogical reasoning work, Idson, Chugh, Bereby-Meyer, Moran, Grosskopf, and Bazerman (2004) suggest that understanding differences, as well as similarities, across problems may also be a very useful means of transferring knowledge. Idson et al. (2004) show that training based on differences can reduce bias in the Acquiring a Company problem, which, as discussed earlier, had proven resistant to many other debiasing techniques. Using the five problems from Tor and Bazerman (2003), Idson et al. (2004) had study participants either (1) examine the two versions of the Monty Hall problem and the two versions of the Dividing a Pie problem as four separate problems, or (2) presented the problems in pairs. All participants were then given multiple trials to solve the Acquiring a Company problem, with pay based on performance. They also gave the same Acquiring a Company problem to other study participants who were not trained on the Monty Hall problem and the Dividing a Pie problem. Idson et al. (2004) found that allowing study participants to view the Monty Hall and Dividing a Pie problems as pairs helped them

understand the differences between the two versions of each problem and generalize the importance of focusing on the decisions of other parties and the rules of the game. These lessons, which were also the keys to solving the Acquiring a Company problem, indeed enabled participants to perform substantially better on this problem. This research offers evidence that examining differences between seemingly related problems may be a successful direction for improving decision making.

What is the optimal level of abstraction that should occur to help people form analogies across problems? Moran, Bereby-Meyer, and Bazerman (2008) argue that teaching people more *general* negotiation principles (such as "Value can be created" or "It is important to understand how parties' interests interrelate") enables successful transfer to a broader range of new negotiation tasks than the focused analogies of Loewenstein et al. (2003). Moran et al. (2008) argue that learning general principles will improve not only the ability to positively transfer specifically learned principles but also the ability to discriminate their appropriateness—i.e., to determine when a principle should and should not be applied.

Moran et al. (2008) found that learners who previously received training in analogical reasoning for one specific negotiation strategy (namely, logrolling issues to create value) did not perform well when confronted with a diverse, face-to-face negotiation with a very different structure. Thus, logrolling may have limited generalizability to other value-creating processes. To test this idea, Moran et al. adapted Thompson et al.'s analogical reasoning training to teach negotiators broad thought processes for creating value in negotiations. Moran et al. (2008) compared *specific training*, wherein learners compare two cases that illustrate the same specific strategy instances (e.g., logrolling), with *diverse training*, wherein learners compare two cases that illustrate different value-creating strategies (e.g., one illustrates logrolling and the other compatibility). Training effectiveness was assessed by looking at performance and outcomes in a negotiation simulation with potential for various value-creating strategies, some of which learners previously had learned and others which they had not.

Moran et al. (2008) found that more diverse analogical training, wherein negotiators learn and compare several different value-creating strategies, fostered greater learning of underlying value-creating negotiation principles than did more specific analogical training. This method facilitated transfer to a very distinctive task and improved performance on a variety of value-creating strategies, including some that participants had never previously encountered. Improved performance was accompanied by a deeper understanding of the potential to create value. Thus, more diverse analogical training can be effective for attaining greater expertise, which fosters understanding of which particular strategies might be effective in different situations and why. At the same time, when training becomes too diverse, the applicability of the message may be lost. The optimal level of abstraction remains an interesting question for future research, as does the question of how analogical reasoning can be applied to improve individual decision making.

STRATEGY 5: TAKE AN OUTSIDER'S VIEW

In Chapter 2, we asked you to estimate ten obscure quantities and to place 98 percent confidence intervals around your estimates. As we noted, most people answer only three to seven of the ten items correctly, despite being 98 percent confident of their intervals. This study bolsters the widespread finding that people are overconfident in their decisions. Interestingly, after people make these ten assessments and are asked to estimate the total number of questions for which the correct answer will be within their confidence interval, these more global estimates are fairly accurate (Gigerenzer, Hoffrage, & Kleinbölting, 1991; Kahneman & Lovallo, 1993). That is, participants generally understand that only three to seven of their 98 percent confidence intervals will actually contain the true estimate!

Kahneman and Lovallo (1993) explain this apparent contradiction by theorizing that we all have two perspectives on decision making: an *insider* view and an *outsider* view. The insider is the biased decision maker who looks at each situation as unique. The outsider, on the other hand, is more capable of generalizing across situations and identifying similarities. Because these two viewpoints exist simultaneously, a member of a consulting team might be well aware that most projects take longer to complete than initial estimates suggest (outsider view) while also believing that her own optimistic estimate of an upcoming project's duration is somehow accurate and unbiased (insider view). Similarly, people who undertake a new home construction or major home renovation know from their friends that such projects typically end up being overdue and 20–50 percent over budget (outsider view). Nevertheless, most people who initiate such a building project believe that theirs will be different—that their home will be completed on time and near the projected costs (insider view).

Kahneman identified a classic situation of insider optimism within a group of colleagues he was working with to define a new curriculum (Kahneman & Lovallo, 1993). The group estimated that the project would take 18–30 months to complete. Kahneman asked a member of the team, who was a distinguished expert in curriculum design, "We are surely not the only team to have tried to develop a curriculum where none existed before. Please try to recall as many cases as you can. Think of them as they were in a stage comparable to ours at present. How long did it take them, from that point, to complete their project?" The team member answered that 40 percent of the projects were never completed, and none were completed in less than seven years. As it turned out, the team took *eight years* to finish its project.

This pattern resonates well with writers. Most of us understand that books take a long time to write; nonetheless, we are optimistic about meeting our own unrealistic deadlines when we sit down to write the first chapter. We may never complete the book, but we will probably believe that the next project will be different. Similarly, Cooper, Woo, and Dunkelberg (1988) found that over 80 percent of entrepreneurs perceived their chances of success to be 70 percent or better, and one-third of them described their success as certain. In contrast, they estimated the mean success rates of businesses similar to their business to be 59

percent. Meanwhile, the five-year survival rate for new businesses is only about 33 percent (Kahneman & Lovallo, 1993).

Kahneman and Lovallo provide convincing evidence that the outsider makes better estimates and decisions than the insider. The outsider view incorporates more relevant data from previous decisions—yet we tend to believe and act on the insider view. Why? Certainly, optimism and overconfidence are factors. In addition, Kahneman and Lovallo document the human tendency to consider all of a decision's various details into our judgment process and, as a consequence, to view each decision as unique. This focus on the here and now leads us to overlook historic data and to let our biases run wild. As a result, we follow the insider view despite the readily available insights of the outsider view.

The insider-outsider distinction suggests another strategy to reduce bias: When making an important decision, invite an outsider to share his or her insight. This may mean conferring with a trusted friend or colleague who has experience with similar decisions. Interestingly, when a friend is building a house, we often predict that construction will cost more and take longer than expected. Our friend is the only one who doesn't know this! So, for decisions that really matter, ask friends you trust for their estimate of what will happen and understand that their outsider perspective may be more accurate than your biased insider view. Alternatively, ask yourself what your outsider self thinks of the situation. To assess this, imagine that the decision was a friend's, and ask yourself what advice you would give him or her. The key is to figure out how to give the outsider a stronger voice in the decision-making process.

STRATEGY 6: UNDERSTAND BIASES IN OTHERS

The nature of managerial life requires you to work closely with the decisions of others, reviewing recommendations, transforming recommendations into decisions, and adjusting decisions made by others in the past. The task of evaluating the decisions of others is fundamentally different from the task of auditing your own decisions. Nonetheless, from reading this book, you have learned that everyone's decisions are influenced to some degree by a shared set of biases. How can you systematically detect bias in the decisions of those around you? Consider the following managerial situation:

> You are the director of marketing for a retail chain that has 40 stores in 14 cities. Annual sales in these stores average between $2 million and $4 million with mean sales of $3 million. Twenty-five of the stores have opened in the last three years, and the company plans to open 30 new stores in the next four years. Because of this growth, you have hired a site location analyst to predict the sales in each potential site. Unfortunately, predicting sales in new markets is difficult, and even the best analyst faces a great deal of uncertainty. As the marketing director, you are evaluated in part by the accuracy of the forecasts coming out of your department. The site location analyst has just given you her latest forecast, $3.8 million in annual sales for a potential site. Demographic data backs up

the analyst's claim that this area should make the store one of the top producers in the chain. What is your reaction to the forecast?

At a naive level, there is reason to have confidence in the analyst's forecast. After all, she knows more than you about the details of the data that underlie the prediction. In addition, your overview of the area also predicts that the store will do well in comparison to existing stores; this evaluation is based on matching the representativeness of this site to other existing sites. The prediction begins to lose force, however, when we consider the prediction in light of a basic but counter-intuitive statistical concept: regression to the mean. In Chapter 3, we saw that the extremeness of our predictions should be moderated toward the mean by the degree of uncertainty in the prediction (Kahneman & Tversky, 1982).

With this rule in mind, let's imagine that the site location analyst is known for her extreme accuracy. In fact, her predictions are almost perfectly accurate and have a correlation of actual sales equal to 1.0. If this is true, it would be appropriate to rely on the $3.8 million prediction. Now let's consider the case in which there is a correlation of zero between the analyst's predictions (based on demographic data) and actual sales. If this is true, her forecast is meaningless, and the only pertinent information is that the average store has sales of $3 million. Therefore, this figure becomes your best estimate. It is most likely, in fact, that the analyst has achieved neither total success nor total failure, but an intermediate level of predictability over the course of her career. The forecast should then fall between sales of the mean store and the analyst's estimate, becoming progressively closer to the analyst's estimate as her ability to predict sales increases (Kahneman & Tversky, 1982). This analysis suggests that, as the director, you will want to reduce the forecast to somewhere between $3 million and $3.8 million, depending on your assessment of the correlation between the analyst's forecasts and actual sales. In essence, the understanding of human judgment taught by this book should help you to systematically adjust the analyst's initial decision.

The preceding analysis offers a rough guide to adjusting the decisions of others. Kahneman and Tversky (1982) have formalized this process into a five-step procedure whose steps are outlined here, using the site location problem as an example. In reviewing each step, you should think about how you might convert this systematic training into an intuitive, natural response. This will allow you, as a manager, to recognize the existence and direction of a wide range of biases across a wide range of decisions and make adjustments accordingly.

1. **Select a comparison group.** This first step consists of selecting the set of past observations to which the current decision or forecast is to be compared. In the site location problem, comparing the new store to the population of all company stores is an obvious group. Other comparison groups often exist. For example, you might decide that only stores that have opened in the last three years are appropriate for comparison, particularly if recent stores are closer in description to the future store than established stores. A more inclusive group allows for a larger basis of

comparison, but its heterogeneity may reduce its comparability to the targeted forecast.

2. **Assess the distribution of the comparison group.** The next step involves assessing the characteristics of the past observations to which the current decision is being compared. If the comparison group consists of all stores, we know the range and mean from the data presented. If we limit the group to recent stores, these data would need to be recalculated. In addition, we might want to get additional data about the shape of the distribution around the mean.
3. **Incorporate intuitive estimation.** This step calls for identification of the decision or forecast of the expert. In this case, the site location analyst's assessment, $3.8 million, is the intuitive estimate that needs to be adjusted. The next two steps attempt to improve this forecast.
4. **Assess the predicted results of the decision.** This is the most difficult step in the corrective procedure, as it requires us to determine the correlation between the decision or forecast and the comparison group data. It may be possible to assess this correlation by comparing past estimates to actual sales. In the absence of these data, you must determine some subjective procedure for this assessment. Kahneman and Tversky (1982) discuss this process in more detail. For our purposes, the key point is that the analyst's estimate assumes a correlation of 1.0 between her prediction and actual sales. In virtually all cases, we must adjust away from this biased estimate.
5. **Adjust the intuitive estimate.** In this step we must calculate the adjustment that reduces the bias error of the initial decision or forecast. For example, this procedure should produce an estimate of $3.8 million when the correlation in Step 4 is 1.0, an estimate of $3 million when the correlation is zero, and estimates proportionally in between when the correlation falls between zero and one. This adjustment can be formalized as follows:

$$\text{adjusted estimate} = \text{group mean} + \text{correlation}\,(\text{initial estimate} - \text{group mean})$$

In our example, it is easy to see that this leads to a prediction of $3.4 million when the correlation is 0.5, $3.6 million when the correlation is 0.75, and so on. The person making the adjustment should fully understand the logic of the procedure and evaluate its relevance to the decision at hand. When arguing for this adjustment, you must recognize that you are likely to face resistance to change.

These five steps provide a clearly delineated process for debiasing an individual's intuition by adjusting for the regression-to-the-mean bias. The formal procedure will typically improve the forecast. More important, a manager who understands the process will become capable of intuitively assessing the degree to which an initial estimate should be regressed to the mean.

This section shows that we can use an understanding of biases to identify systematic error in the decisions of others. Adjusting for regression to the mean is simply one example of how such a technique can be systematized. When we consult with organizations, our knowledge of the various biases documented in this book allows us to identify biases across a variety of problem types.

We now have a model for adjusting a wide range of biased decisions in both individual and multiparty contexts. Broadly, it involves three phases. First, we need to accurately perceive and analyze the context within which the decision is being made. Next, we need to distinguish the potential bias(es) surrounding the decision and the decision makers. Finally, we need to identify and make the appropriate logical adjustments for that decision. This judgment-improvement technique can be used to evaluate and adjust our own, as well as others', intuitive judgments in a variety of situations.

You can also use your new knowledge of the biases of others to identify optimal moves in a competitive environment. Richard Thaler, whose ideas we have cited often in this book, teamed up with Russell Fuller to create the Fuller-Thaler mutual funds (www.fullerthaler.com). These funds buy securities by taking advantage of the predictable biases of key market participants. Fuller and Thaler argue that these biases result in mispricing of securities. For example, they argue that most analysts underreact to new, positive information about firms. By identifying how decision biases create under- and overvalued firms, Fuller and Thaler have created funds that outperform the market.

STRATEGY 7: NUDGE WISER AND MORE ETHICAL DECISIONS

Which option do you prefer (from Bazerman, Baron, and Shonk, 2001):

(a) If you die in an accident, your heart will be used to save another person's life. In addition, if you ever need a heart transplant, there will be a 90 percent chance that you will get a heart.

(b) If you die in an accident, you will be buried with your heart in your body. In addition, if you ever need a heart transplant, there will be a 45 percent chance that you will get a heart.

In this problem, most people chose (a). So why does the United States maintain an organ donation policy that resembles (b)? The answer lies in the psychology of the evaluation of losses and gains. As we discussed in Chapter 5, Tversky and Kahneman (1991) have documented that losses loom larger in our minds than gains. Moving to an opt-out program would save lives (an important gain) but would also have costs salient to some individuals, such as the prospect of being buried without all of their organs.

As a result, in the United States alone, about 50,000 people are on waiting lists for organs at any given time. More than a third of them will die before an

organ is found. The number of organ donors has declined in recent decades, due to increased use of seatbelts and motorcycle helmets, and only 4,500 of the 11,000 eligible donors actually donate their organs. If we could double this figure, we could save an additional one-quarter of the approximately 15,000 people who die each year in the United States because of the lack of organs.

This situation exists despite the fact that we know how to increase the number of organs available for donation. Bazerman et al. (2001) argued that like many other countries (including Austria, Belgium, France, and Sweden), we could presume consent to organ donation (an opt-out program) rather than presuming non-consent (an opt-in program). That is, we could change the default in the United States to assume that eligible people are organ donors upon death unless they specifically opt out of the organ-donation system. Thanks to the clever empirical work of Johnson and Goldstein (2003), we already know what the result would be. European countries with an opt-in program similar to that of the United States have organ donations rates between 4 and 28 percent. In contrast, European countries with opt-out programs have rates ranging from 86 to 100 percent.

Enormously costly inefficiencies such as the U.S. organ donation system are surprisingly common in society. In their fascinating book *Nudge*, Thaler and Sunstein (2008), outline a structure for thinking about devising more efficient and beneficial organizational and societal systems. They argue that we can anticipate the mistakes humans make on a regular basis and then create systems that correct for these mistakes in a way that will nudge them toward better and more ethical decisions. Thaler and Sunstein's "libertarian paternalism" is libertarian in the sense that people have control over maintaining or expanding the options available to them and paternalistic in the sense that the system's architects attempt to guide people toward wiser decisions.

One of the most famous examples of Thaler's prior work in the area of nudges comes from a study by Thaler and Benartzi (2004) focusing on how to increase employees' enrollment in retirement plans that will benefit them over the long term. Using the psychological principles described in our book, they motivate people to increase their contributions to their 401(k) plans through a program called "Save More Tomorrow." Under this program, workers have the option of committing in advance to increase their retirement savings rates when they get a raise. The creation of the program was based on an understanding of the concepts of discounting, procrastination, and loss aversion. The design encourages commitment because people are more likely to choose what they know they *should* do when considering future rather than present events. The program remains effective over time thanks to inertia: people rarely take the initiative to opt out of the program once they have committed to it. Finally, the contribution increases are not difficult for the saver to stomach because the savings rate increases with the size of one's paycheck and never leads to a decrease in disposable income. The additional savings come from foregone consumption of anticipated gains rather than from a decrease in current disposable income.

In just over two years, the Save More Tomorrow pilot plan more than tripled the savings rates of those who joined. Since then, numerous retirement-plan administrators have implemented the Save More Tomorrow concept, including Vanguard, T. Rowe Price, TIAA-CREF, Fidelity, and Hewitt Associates. According to the Profit Sharing Council of America, as of 2007, 39 percent of large U.S. employers have adopted an automatic retirement contribution escalation plan (Thaler & Sunstein, 2008). Automatic enrollment dramatically increases participation in such programs. The Safelite Group, the first to adopt an "opt out" enrollment (in which employees must actively decline participation), automatically enrolled 93 percent of program participants in 2003. Only 6 percent chose to opt out over the next year, leaving the bulk of participants to save much more than they would have if they had been required to actively "opt in" to the program. Thaler and Sunstein (2008) provide numerous examples of common-sense nudges that suggest how, by thinking about human barriers to wise decisions, we can design systems that will lead to more positive results.

Nudges can be quite simple. Bohnet, van Geen, and Bazerman (2012) focus on how to nudge employers to make personnel decisions based on individual capabilities rather than stereotypes. They start by showing that when evaluating employees one at a time, many people rely on gender stereotypes: they select men for mathematical tasks and women for verbal tasks. When the hiring system is adjusted so that two or more potential employees are considered jointly, the focus of decision makers shifts to the ability of the potential employees, and they make decisions that are more ethical to the job candidates and that lead to better organizational performance. By making such small changes in how we make common decisions in organizations, we can inspire wiser and more ethical decisions.

CONCLUSION

In this final chapter, we have introduced seven strategies for correcting the deficiencies in our decision making. The first three strategies seek to create broad change in our intuitive responses to decision-making situations. In general, they strive to heighten our awareness of our cognitive limitations and our susceptibility to bias. The last four strategies provide techniques for improving specific decisions in specific contexts. They offer concrete methods for testing and adjusting actual decisions. Together, these seven strategies provide tools for changing and "refreezing" your intuitive decision-making processes in the future.

An optimistic but naive view of this book is that you are now immediately capable of improving your decision making. Why naive? Because it would be premature to expect you to have fully integrated the process of changing your judgment for the better. If unfreezing did not take place, then the book failed. If you were not provided with sufficient information for change, the book again failed. However, the responsibility for refreezing new processes and using the decision-improvement strategies suggested in this last chapter lies with you.

Refreezing requires a period in which you constantly review your decision-making processes for the errors identified in this book. Refreezing also requires that you be vigilant in your search for biases in the more complex world of decisions that you face. Creating lasting internal improvement in decision making is a complex task that occurs gradually over time through persistent monitoring. It is far easier to identify a bias while you are reading a book about decision making than when you are in the midst of an organizational crisis. Raiffa (1984) has found that his students were likely to use appropriate decision-making strategies on one of his exams but failed to generalize the relevance of these strategies to similar problems in courses taught by other instructors. Thus, making adjustments to your decision-making processes requires constant attention.

In addition to improving your own decisions, the ideas in this book should be very useful for informing you about the decisions of others. We are often faced with situations in which we are suspicious of another party's decision making, but we lack the vocabulary to articulate the flaws in their logic. This book offers systematic clues for understanding and explaining the biases of others. You can practice spotting others' biases while reading the newspaper or watching a sporting event on television. Reporters, sportscasters, politicians, and other information providers and public servants constantly make statements that exemplify the biased decision-making processes outlined in this book.

We hope that this book has dispelled some of your assumptions about decision making. We also hope to have raised your awareness of the importance of the decision-making process itself, rather than just the results of this process. We are disturbed by the fact that most managers reward results rather than good decisions. As we have seen, managers make many decisions for the wrong reasons. Nevertheless, because so many important decisions involve uncertainty, plenty of good decisions turn out badly, and some bad decisions turn out well. To the extent that a manager rewards results and not sound decision making, the manager is likely to be rewarding behaviors that may not work in the future.

Davis (1971) argues that "interesting" writing leads readers to question issues that they never thought about before. Thus, identifying new issues may be more important than providing new answers to old questions. In this sense, we hope this book has succeeded at being interesting by making you aware of aspects of your decision-making process that inspire new questions and solutions.