

(Topping, 1998). In the medical realm, peer assessment, coupled with self-assessment exercises, may prove a significant help with health care training in developing countries where formal supervision is expensive to maintain. In the late 1990s, the Indonesian government, in concert with several international organizations, conducted a training program with local health-clinic counselors to improve their communication skills when discussing contraception with clients. Counselors attended a 5-day refresher seminar and then were assigned to one of three groups. For the next 16 weeks, one group completed self-assessment exercises on their communication skills. The second group completed self-assessment exercises, plus attended weekly peer-review sessions. The third group served as a control and did not receive any intervention.

Six months after training, researchers taped and coded two interviews involving each counselor, to assess the extent to which each counselor provided medical information and exhibited a facilitative and active communication style. Both interventions increased facilitative communication—with the intervention involving peer assessment producing more improvement in how much information (e.g., about lifestyle consequences) clients were given. As a consequence, clients talked more actively with counselors who had received the intervention that combined self- and peer assessment (Kim, Putjuk, Kols, & Basuki, 2000). The costs associated with these self- and peer-assessment interventions turned out to be low given the effects these interventions brought about.

Peer assessment, however, is not without difficulties. It can be demanding and time-intensive. It can provoke anxiety (Topping, 1998), as well as fears that grades will be more a product of illegitimate influences (such as friendship) than of objective attainment (Dancer & Dancer, 1992). Poor students do not provide as accurate assessments as their more skilled peers do (Saavedra & Kwun, 1993). Thus, using peer assessment is not guaranteed to improve student performance unless the assessments are conducted carefully. Peer assessments become more valid as they are based on a larger number of observations and a greater number of dimensions of skill. They are also most helpful when standards are clear and more than one peer provides an assessment. Peer-assessment exercises are also enhanced if instructors communicate the purpose of the exercises clearly, articulate the dimensions of judgment clearly, provide training when necessary, and monitor students' evaluations, intervening when they are too harsh or too lenient (Norcini, 2003). Peer assessments better match teachers' evaluations when students provide global evaluations based on a number of well-defined criteria rather than a number of separate evaluations along several individual dimensions (Falchikov & Goldfinch, 2000).

## THE WORKPLACE

Given the research reviewed earlier in this monograph, it is not surprising that biased self-views also affect success and social

relationships in the workplace. When one looks at the accuracy of self-assessment in the workplace, from the office cubicle to the executive boardroom, one sees that people tend to hold overly inflated self-views that are modestly related to actual performance. One also sees reverberations of other themes (overconfidence in judgment, egocentric neglect) that are apparent in the domains of health and education.

If people often have biased self-views, it is not surprising that these biases affect workplace outcomes ranging from the ability to anticipate performance evaluations to the ability to gauge promotion opportunities. What is perhaps less apparent is that organizations can adopt many routines and procedures that keep biased self-views from affecting performance in many situations. As a result, biased self-views may have their most systematic and damaging effects in the circumstances in which individuals are least likely to face routines and procedures that could repair their biased decisions, for example, when people confront new projects that have few precedents, or when CEOs and management teams at the very tops of organizations, where there are few routine correctives, make important, large-stakes decisions about which new markets to enter and which companies to acquire.

### Accuracy of Self-Knowledge in Organizations

In organizational life, the largest surprises generated by lack of self-knowledge may be those that are produced when self-evaluations are not echoed by supervisors who set raises and hand out promotions. If employees overrate their own performance, it is difficult to imagine how people could not wind up disappointed at least some of the time. For example, Zenger (1992) studied several hundred engineers at two high-tech companies and found that 32% of the engineers in one company and 42% in the other rated their own performance in the top 5% of all engineers. Imagine the difficulty of conducting honest performance evaluations for these engineers.

Although workers may find it tempting to blame their supervisors when they receive a less-than-stellar evaluation, research suggests that they should trust their supervisors' views more than their own. Typically, the views of other people—subordinates, peers, and superiors—agree with each other more often than with self-views. In one typical study, Bass and Yammarino (1991) studied U.S. Navy officers who had graduated from the U.S. Naval Academy and were on active duty on ships ranging from tenders to aircraft carriers. The researchers collected leadership ratings from the officers' subordinates and the officers themselves. They also retrieved some measures of leadership from the officers' superiors—each officer's performance on regular "fitness reports" and promotions relative to peers. On all seven positive leadership dimensions studied (e.g., charisma, individual consideration, intellectual stimulation), the officers rated themselves more positively than did their subordinates. But although the officers' self-ratings displayed almost no correlation with

their own fitness reports and promotions, the subordinates' ratings inevitably did.

This is one example, but its message is echoed across many studies in many kinds of organizations. Harris and Schaubroeck (1988) conducted a meta-analysis of 30 years of academic research comparing the self-ratings of individual employees with the ratings of their supervisors or peers. Although the ratings of peers and supervisors agreed pretty well ( $r = .62$ ), their correlation with people's own ratings of their job performance was lower ( $r = .35$  for supervisors' ratings,  $r = .36$  for peers' ratings). This lack of agreement between self-ratings and outside ratings was exacerbated when there was more ambiguity: when the rating scale was ill-defined (e.g., correlations were lower when people assigned a single global rating rather than separate ratings on specific task dimensions, and they were lower when people rated abstract traits rather than specific behaviors) or when job descriptions were ambiguous (correlations were lower for less-well-defined managerial or professional jobs than for blue-collar or service jobs).

For a number of reasons, individuals may find it difficult to learn about how well they are performing (Ashford, 1989). First, the environment may be tricky. Organizational standards of performance may be shifting, complex, or subject to disagreement. Second, individuals may avoid seeking feedback—getting feedback requires them to balance the potential benefits of improvement against the cost of learning something that would be a blow to self-esteem. Also, individuals may resist seeking feedback even if they might like to have it because asking for feedback might make them appear insecure or needy.

### Challenges in Providing Employees With Feedback

If people are not very good at assessing their own performance, and if they are reluctant to solicit feedback on their own, cannot organizations adopt routines (e.g., regular performance reviews or evaluations) that could clue employees in to their strengths and weaknesses? Although the organizational literature remains perpetually optimistic about solving this problem by encouraging regular and reliable feedback, actual feedback systems tend to be relatively ineffective for the following reasons:


- Feedback is *infrequent*: A study by the American Management Association in 1984 found that 75% of organizations provided formal feedback only once a year (Ashford, 1989). When feedback is this infrequent, it takes on additional freight and significance, exacerbating the other problems in this list.
- Feedback is *threatening*: Because people are overconfident about their abilities, true feedback will often be perceived as negative and, potentially, threatening. Feedback may not work if people experience it as commentary about their character or future outcomes (e.g., “Is she saying I’ve been a bad manager?” “Will I lose my bonus this year?”) rather than about the way they are approaching a specific task (e.g., “Let’s think of ways that your team could produce 40 more widgets a week”); see DeNisi & Kluger, 2000; Kluger & DeNisi, 1996).

- Feedback is *sugarcoated*: People generally find it painful to receive negative feedback, so evaluators find it painful to give negative feedback. Thus, evaluators may withhold negative feedback or distort it in a positive direction.
- Feedback is *given too late*: Because evaluators are hesitant to disillusion the people they need to evaluate, they may wait so long to say anything that by the time they do, they are beside themselves with anger (Larson, 1989).

Various solutions for making feedback more effective have been suggested and explored. Some researchers have argued that feedback should be separated from evaluation, and that feedback should be provided more frequently (e.g., “coaching”); but in a busy organization where formal reviews are done at most once a year, asking for less formal feedback multiple times a year might be unrealistic. Other researchers have advocated taking the burden of reviews off one individual manager (e.g., by using 360° reviews, in which a focal individual may be evaluated by peers, subordinates, customers, and multiple superiors), but this makes the evaluation process more costly and intense. Some organizations have tried to force managers to rate their employees along a curve—only a small percentage of people can be rated as high performers, and half of employees must be rated below average—but such systems frequently produce frustration and controversy (Baron & Kreps, 1999, chap. 10).

Also, because of the basic problems of self-assessment, any of these proposed solutions is subject to its own potential problems. Evaluators may give feedback infrequently because they find it emotionally taxing to give people feedback that may violate their self-image. The burden of feedback systems falls especially hard on front-line managers, who have to give feedback that disappoints or angers employees who assume they are above average and cannot see why their managers do not agree. Giving people feedback more frequently, although it may reduce the stakes and allow for more specific instruction that would be less threatening, forces evaluators to repeat a painful process more often. Because of these problems with giving feedback, many organizations evolve to an equilibrium pattern of feedback that may be sub-optimal for organizational performance, but more viable from the standpoint of interpersonal relationships.

Organizations must select the best people for further development and promotion, but as we have indicated in our discussion of evaluation and feedback, this selection process is likely to be contentious when a substantial majority of workers think they are above average. How, then, can organizations recognize and develop their best employees without undermining the motivation of the bulk of employees in the middle? One way might be to give high and low evaluations to a few exceptionally good and bad performers, respectively, and to give the bulk of the employees ambiguously positive evaluations. Figure 1 shows an actual distribution of evaluations for a high-technology company studied by Zenger (1992). Note that relatively few people (around 10%) were given the very top rating, but even fewer (less than 5%)



**Fig. 1.** Performance evaluations at a high-technology company (Zenger, 1992). The graph shows the percentage of employees given each rating on a scale from 1 (top performance category) to 5 (bottom category).

were given ratings (4 or 5) that were explicitly below average. The majority of employees were clustered in a large category that, at least in terms of the overt scale, rated them solidly “above average.” Here, as in Garrison Keillor’s Lake Wobegon, everyone can be above average; but this organization gives people evaluations that prove it.

Zenger (1992) noted that such rating systems may reduce employees’ concerns with the evaluation system. When any particular employee is singled out for special positive recognition, this makes many other employees—all those who correctly or incorrectly feel they are performing better—feel that their efforts have not been equitably rewarded. Organizations with positively skewed evaluation systems may get away with singling out a few really high performers—those who peers agree are clearly deserving—if they avoid tougher distinctions among people in the middle of the distribution, where quality distinctions may be more contested. Such evaluation systems cleverly allow organizations to select the best (and get rid of the worst) with the fewest ruffled feathers, but the problem is that these systems do not provide people in the middle of the distribution with an evaluation that may signal them to improve their performance. The overall results of such systems may depend on whether organizations can separate feedback from evaluation and provide feedback that encourages people in the middle to improve even though they are already being evaluated as “above average.”

#### Cognitive Repairs for Mistaken Self-Judgment

For people at the bottom of a hierarchical organization, lack of self-knowledge may hinder career or personal success. This is a recipe for personal unhappiness, but it may not affect how they complete their day-to-day tasks. Luckily, organizations often develop or evolve reminders, routines, and procedures that help mitigate problems that could be caused by employees who lack awareness of their own abilities: Marketing managers have their ad campaigns checked by higher-level managers who have created campaigns before, and loan officers in banks must have their loans approved by loan audit committees. Combined, these

standardized organizational routines and cross-checks may prevent employees from making big mistakes through their own lack of self-knowledge.

#### *Adding Safety Factors and Buffer Time*

For example, engineers are often overconfident about their ability to calculate the amount of concrete needed for a dam, or the strength of materials needed for an airplane wing. This lack of self-knowledge could have obvious negative implications for people’s safety, but the engineering profession has evolved a “repair” for this overconfidence: safety factors (Heath, Larrick, & Klayman, 1998). Thus, after engineers calculate the amount of concrete they need to hold back a given amount of water for a dam, they multiply this number by a safety factor, a number between 3 and 8, thereby tripling, or even octupling, the amount of concrete their calculations suggest. Are engineers overconfident? Yes. Are they overconfident by a factor of 8? No.

Similar repairs have evolved in other areas. The need for repairs of overconfident estimates in, say, software development is evident DeMarco (1982), a respected consultant on software development, defined an estimated completion time as “the most optimistic prediction that has a non-zero probability of coming true” (p. 14). Software developers at Microsoft often experience burnout because they “grossly underestimate” how long it will take them to accomplish particular tasks (Cusumano & Selby, 1995, p. 94). This hurts the programmers who struggle to accomplish overly difficult goals, and it hurts the organization that finds itself unable to deliver the products it promised at the time it promised them. Microsoft repairs this problem by adding buffer time to projects. In planning the timeline for developing stand-alone applications such as Excel and Word, Microsoft builds in buffer time equal to 30% of the schedule. For more complex projects, such as operating systems, it might add buffer time of 50%.

Microsoft has also developed processes that help individual programmers assess their own self-knowledge more correctly. Said one manager quoted by Cusumano and Selby (1995), “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' " (p. 254). Liberman and Trope (1998) showed that when people consider tasks that are distant in time, they analyze them at a higher level of abstraction than they do tasks that are close in time; the repair for this problem at Microsoft—breaking an overall task down into 22 subtasks—is a textbook application of Trope and Liberman's (2003) research, but it developed as an organizational practice long before researchers had recognized and elucidated the psychological mechanism.

#### *Forcing People to Pay Attention to the Environment*

Other repairs help people pay attention to the situational forces that affect their outcomes. People habitually take credit for successes, whether caused by their own actions or by situational luck (Mullen & Riordan, 1988), and so do traders on Wall Street who are making particularly high-stakes decisions. Traders have evolved a simple repair to force each other to pay attention to the broader market. They tell each other, "Don't confuse brains and a bull market" (Odean, 1998).

#### *Providing Feedback From Other People*

Often, overly optimistic self-assessments can be repaired by forcing optimistic individuals to answer questions from other people who are not as personally committed to a project and are therefore better able to analyze it objectively. The norms of science emphasize rigorous, thoughtful attention to detail, but even well-socialized scientists cannot always recognize the limits of their research. One study that compared successful and unsuccessful microbiology labs found that successful labs, where many scientific breakthroughs occurred, placed more emphasis on group lab meetings in which researchers were forced to answer the questions of their interested but skeptical peers from diverse areas of research (Dunbar, 1995). Disney holds the "Gong Show" for new ideas, during which (as in the 1970s game show) participants bang a gong at a point where a pitch is becoming unconvincing, and the Pentagon requires that important missions be vetted by the "Murder Board," a group of senior combat officers who will kill a mission that is not adequately planned (Heath et al., 1998).

#### **The Problem of the New**

The fact that many organizations have evolved cognitive repairs to correct the distorted self-views of their employees suggests that inadequate self-knowledge creates problems for organizations, and that those problems are important enough for some organizations to have invested the effort to design useful correctives. But not all problems are equally likely to be repaired, and current

research does not indicate how often problems are effectively repaired. On the basis of research on individual and organizational learning, we speculate that repairs are most likely to develop when tasks are repeated, when feedback is relatively prompt, and when other employees are in a position to alert the focal decision maker to his or her own lack of awareness (Heath et al., 1998). Interestingly, this suggests that problems with self-knowledge are likely to be particularly pronounced when organizations undertake novel tasks and when top managers, who do not face strong social feedback from peers, are the ones taking action.

#### *Novel Projects*

Organizations are unlikely to be able to correct for the inadequate self-knowledge of their employees when they engage in a task for the first time. Existing repairs are unlikely to be effective for truly novel tasks, and the organizational literature is filled with anecdotes about novel, large-scale projects that provide evidence of the planning fallacy writ large. For example, the Sydney Opera House embodied a novel design that stretched the architectural and engineering sophistication of its day. In 1957, the building was predicted to open in 1963 at a cost of \$7 million. After the project was scaled down, the opera house eventually opened a decade later, in 1973, at a cost of \$102 million (Buehler et al., 2002).

But such planning disasters are not limited to single isolated cases. When researchers have systematically examined arenas in which organizations undertake novel projects, they have found that such projects frequently go very wrong. One study of pioneer process plants (i.e., the first plants built to produce chemicals using a new chemical process) showed that the typical plant experienced actual construction costs that were double the original estimate. A study of start-ups showed that more than 80% fell short of their projected market share (Davis, 1985). The U.S. General Accounting Office estimated that when the military purchased equipment involving new technology, it was delivered on time and within budget only 1% of the time (Buehler et al., 2002).

Such broad-scale planning disasters could conceivably be created by a number of social, institutional, and economic factors, but research demonstrates that lack of self-insight, and the planning fallacy specifically, can cause such surprises even when completing a novel project is largely under the control of one individual. One cultural influence that potentially exacerbates the planning fallacy, particularly in the case of novel projects, is socially shared views of management. According to a classic definition, managers "plan, organize, direct, and control" (see Mintzberg, 1975). This definition focuses heavily on the inside view; it encourages managers to picture themselves in charge, directing and controlling what happens. But an important part of self-insight is realizing what one does not control. One good antidote to excessive optimism is realizing, "It's not about you." When people focus on themselves and what they intend to

do, they typically do not consider the large number of ways the environment can react to hinder the best-laid plans.

Not only may definitions of management exacerbate problems with self-knowledge, but the frameworks people use to assess situations may do the same. For example, the field of business strategy provides advice to top managers about how to make major decisions about which new markets to enter. One type of strategic analysis, industry analysis, focuses on the characteristics of various industries that may make them profitable or not (e.g., the power of suppliers or customers, threats from current competitors or potential ones; e.g., Porter, 1980), whereas other types focus more on the internal skills and capacities of the firm (e.g., the resource-based view, discussed by Barney, 1991, or core competencies, discussed by Prahalad & Hamel, 1990). Although both types of strategic analysis are important, psychological research suggests that analysis focusing on the characteristics of industries is likely to be a more effective repair to the initial assessment tendencies of top managers. Left alone, they are more likely to take an inside view, considering the core competencies or resources of their firm, than they are to take an outside view and consider the situational dynamics and history of their industry.

#### *Entering New Markets*

If lack of self-knowledge is most damaging when tasks are novel and when there are few social corrections for mistakes, then, ironically, the major negative impact of limited self-knowledge will happen at the upper levels of organizations. Organizations that have elaborate procedures and committees to avoid making a bad \$10,000 loan may have no routines or social correctives to prevent their CEOs from making a bad \$100,000,000 decision to acquire another firm or enter a previously untested market.

Most new businesses fail quickly—61.5% exit within 5 years, and almost 80% within 10 years. Why would entrepreneurs enter markets when their businesses have such a small chance of success? One possible reason is that good business opportunities do not last very long. Businesses that “fail” may actually be rational gambles that take advantage of a short-lived market opportunity. Or perhaps entrepreneurs enjoy taking risks; they may be willing to gamble on a small probability of a fabulous success. A third possibility is that lack of self-awareness plays an important role in decisions to enter new markets: Perhaps entrepreneurs are aware that the average firm is unprofitable, but think that they have unique skills that will allow them to succeed even if most other people fail.

In a series of clever experiments, Camerer and Lovallo (1999) provided evidence that supports this last possibility. They used laboratory-based economic games that participants played for substantial sums of money. The games forced participants to make decisions about entering competitive economic markets; a small number of those who entered a market would earn profits based on their competitive “rank,” but everyone else would lose money. In one market with 8 participants, for example, the

highest-ranked competitor would win \$33, the next-highest would win \$17, and everyone else would lose \$10.

The key experimental manipulation was whether competitive rank was determined randomly (by drawing a number) or through skill (at, say, solving logic problems or answering trivia questions). On the basis of previous research on lack of self-knowledge, Camerer and Lovallo (1999) predicted that people would enter the markets too often when competitive rank depended on skill. Overall, in economic markets where rank was determined randomly, the typical participant was profitable 77% of the time, but when rank was determined by skill—and participants could tell themselves why their own abilities might make them better than others—the typical participant was profitable only 40% of the time.

This effect was magnified when participants decided to participate in the experiment after being told up front that the outcomes in the experiment would be based on skill. Some participants signed up for “an economics experiment” without learning anything else about it; others signed up for an experiment in which their earnings would be based on their skill at answering sports or current-events trivia, and they were told that “people who are very good might earn a considerable sum of money.” Participants in the latter experiments were reminded during their experiment that all the other people in the room had also nominated themselves on the basis of their own expertise, but these self-selected participants performed much worse than the participants who simply signed up for an economics experiment: Non-self-selected participants were profitable in 71% of the skill rounds, but the self-selected participants ended up profitable in only 6% of the skill rounds, thus losing money 94% of the time.

These experimental results serve to explain some behavioral patterns among real businesses. Although most new businesses fail, the failures are not evenly distributed. Moore and Cain (2004) noted that some industries—such as retail clothing stores, restaurants, and bars—are marked by persistently high rates of entry and exit. Although people typically rate themselves above average, this tendency is reduced or reversed in domains that are commonly regarded as difficult (e.g., juggling or chess; Kruger, 1999). As we explained earlier, people tend to focus myopically on their own abilities when making comparisons, so they think they will succeed in easy domains and fail in difficult domains, regardless of the skills and abilities of their competitors.

Moore and Cain (2004) argued that in industries believed to be easy, many entrepreneurs will think they have what it takes to succeed. They studied this hypothesis in a laboratory game, similar to the market-entry game just described, in which ranking was based on performance answering difficult questions (e.g., “What is the closest star outside our solar system?” Answer: Proxima Centauri), performance answering easy questions (e.g., “What is the common name for the star inside our solar system?” Answer: the Sun), or a randomly generated number. On average, too many people entered the economic markets when ranking was based on answers to simple questions, and as a result the typical

participant lost money. However, when ranking was determined by answers to difficult questions or by random drawing, fewer people entered the markets, and the typical participant made money. Interestingly, the losses on the simple rounds occurred even though people predicted that many others would also enter—they just thought they would do better than others at the simple task. This suggests that opportunities for success may be more common in difficult domains than in easy domains. Easy domains attract too many people, whereas hard domains scare people away. In both cases, people would do better if they thought seriously about their potential competitors and not just their own skills and abilities.

Landier and Thesmar (2003) studied the real-world impact of self-assessments in a sample of 23,000 French entrepreneurs. They measured overconfidence by comparing how entrepreneurs forecasted the growth of their firms with the firms' actual results several years later. If the entrepreneurs overestimated the future growth of their firms, they were classified as overconfident. Entrepreneurs were more likely to be overconfident when they were working to implement a new idea rather than a standard business venture (knowledge of the situation is likely to be worse in the case of a new idea) and when they expressed a greater desire for personal autonomy (a result that perhaps reflects a motivational bias).

Landier and Thesmar (2003) found that overconfidence had some positive effects. For example, during the initial stages of their businesses, entrepreneurs who were overconfident worked harder in looking for potential customers and in researching technology, relative to those who were not so confident. But overconfidence also had disadvantages. Entrepreneurs who were overconfident were also more likely to finance their projects with the money they raised informally from themselves, friends, and family as opposed to a bank loan, but when they got a bank loan, they were more likely to rely on short-term (as opposed to long-term) borrowing. Both of these financing options gave the entrepreneurs more control in the short term, but they also exposed the entrepreneurs to greater long-term risk. When problems occurred, overconfident entrepreneurs did not abandon projects quickly enough, and many ended up losing money. In contrast, other entrepreneurs, especially those who had borrowed money from banks, were more likely to terminate their project in time to avoid losing their own money.

In general, psychologists have assumed that optimism and overconfidence are fairly general phenomena that most people share. But by ignoring variation in optimism or overconfidence, psychologists miss the fact that when situations allow people to decide whether to participate or not, this self-selection may ensure that the people who decide to participate are the most optimistic. When someone is among the most optimistic of many contenders, this person's optimism need not be a personality variable, relatively constant across situations; it may simply indicate that the individual has made a mistake in analyzing a specific situation. Interestingly, among Landier and Thesmar's

(2003) French entrepreneurs, education increased overconfidence. In general, one might think that education would provide knowledge that would reduce overconfidence, but because educated people have good options other than starting a new business, the only highly educated people who will give up their existing, already-attractive jobs will be those who are most confident.

Self-selection is particularly important in understanding how common overconfidence will be in social situations in general. Researchers have shown that optimistic self-assessments can be reversed with tasks that are sufficiently difficult. So predicting whether people will be more often optimistic or pessimistic comes down to a question about the environment—is the world filled with tasks that are perceived as easy (which prompt optimism) or tasks that are perceived as difficult (which prompt pessimism)? If environments typically allow self-selection, optimism is probably more likely than pessimism. Even if the world is filled with tasks that are, in reality, hard, people are likely to pursue only the ones they think look easy. Many people may not believe they can sing opera or win programming contests, but many may think they could run a pretty good restaurant or bar. In the laboratory, psychologists can randomly assign people to tasks, even hard ones, but in the world, people self-select situations that seem to fit their self-assessed skills and abilities. Self-selection becomes particularly important when people choose to compete with others who also think they are quite capable. Egocentric neglect becomes more important when people face competitors in a more highly selected group. Senior managers have survived many rounds of a "promotion tournament," and CEOs have survived them all.

### The Overconfident CEO

CEOs illustrate the problems of overconfidence in the starkest form. They are highly self-selected, so they are particularly prone to overconfidence. And because they sit at the top of their firms, they have fewer checks on their decisions and face fewer organizational repairs, so any overconfident decisions they make may go uncorrected. In this section, we consider the role of overconfidence in CEOs' decisions about how to invest free cash flows and when to acquire other companies.

#### *Investing Free Cash Flow*

In one study, Malmendier and Tate (in press) measured overconfidence by looking for CEOs who were heavily invested in the stock of their companies and who resisted opportunities to sell. Economic theories would predict that such CEOs make the wisest decisions for their companies because their own money is heavily tied up in their firms. But Malmendier and Tate argued that these CEOs are also implicitly expressing high confidence—perhaps overconfidence—in their own ability to manage their firms.

Research in corporate finance has identified an interesting bind faced by CEOs who are overconfident about their own

abilities. On the one hand, they feel great confidence in their ability to start projects that will create economic value, but on the other hand, they find it difficult to raise the money to start those new projects. Financial markets may not share their overly optimistic opinions of their prospects and abilities, and may make it difficult for them to raise money. The stock market may place a low price on their stock, making it hard for them to sell more stock to raise money. Or banks may impose a high interest rate on bank loans to compensate investors for the risk that overly optimistic projects might go into default and not make enough to repay the loans. Thus, overconfident CEOs see a world filled with promising opportunities, but they find it frustratingly difficult to raise money to pursue these opportunities because the financial markets do not share their optimism.

One way CEOs can escape this bind is by starting new projects using only the cash that their organizations generate internally (Heaton, 2002; Malmendier & Tate, in press). Some firms generate more *free cash flow* (i.e., cash that is not committed to paying off loans, etc.) than others, and overconfident CEOs who run such firms have more cash available to devote to “promising” projects.

Malmendier and Tate (in press) found that CEOs who are heavily invested in the stock of their companies are the ones most sensitive to free cash flows—they invest more in new projects when internal cash flows are high and less when internal cash flows are low. By using internal cash to finance their projects, these CEOs avoid having to get outside commitments from financial markets that may not agree with their optimistic views. Interestingly, CEOs are also more likely to invest free cash flows when they hold two other titles, president and chairman of the board. When CEOs are allowed to accumulate all these titles, this may be an indicator that they face weak oversight by their boards of directors. In sum, overconfidence may cause CEOs to see projects as promising even though other people think those projects are risky. And they are more likely to pursue those projects when they can do so without a second opinion from the external financial markets or an independent, active board of directors.

### *Mergers and Acquisitions*

Acquisitions, in which the management of one firm pays to take control of another firm, are a particularly interesting test of self-knowledge (Malmendier & Tate, 2003). Implicitly, the management of the acquiring firm is claiming that they can manage the acquired firm *better* than its current management does. Moreover, they often pay a substantial premium to take over the other firm, and this premium serves as a proxy for how much better they think they can do. Between 1976 and 1990, such takeover premiums averaged 41%, and many were over 100% (Hayward & Hambrick, 1997). These figures are a substantial statement of self-confidence, roughly equivalent to the managers of the acquiring firms claiming that they can manage the acquired firms 41% better than the existing managers. This level of confidence is rarely echoed by external observers. Indeed, on the day most

mergers are announced, the combined stock price of the two firms involved generally *falls*, indicating that the market—composed of people who are voting with their money—predicts that on balance the combined firm will be less healthy than the two firms would be separately (Hayward & Hambrick, 1997; Malmendier & Tate, 2003).

Hayward and Hambrick (1997) investigated the idea that corporate acquisitions might be driven by managerial hubris (Roll, 1986), using a sample of acquisitions that were each valued at over \$100 million. They measured three variables as proxies for hubris. The first was recent organizational success. Hayward and Hambrick argued that such success might foster hubris because people often take credit for their success even when it is caused by luck or general situational conditions. Leaders, in particular, get disproportionate credit for successes (Meindl, Ehrlich, & Dukerich, 1985). The second proxy was media praise, which might accentuate hubris because “romantic” media portrayals might foster the CEO’s self-impression that he or she is clever and in control. Media portrayals might also crystallize and cement the CEO’s power inside the firm as other people come to believe in his or her abilities. The third proxy was relative power within the firm, as measured by the ratio of the CEO’s compensation to that of the next most highly paid corporate officer. CEOs typically receive 30 to 50% more than their closest colleagues, but are often paid more than twice as much. A relatively high ratio means that the CEO’s hubris is unlikely to be checked by social interaction with “inferiors.”

Hayward and Hambrick (1997) discovered that all three measures of hubris independently predicted increased premiums for acquisitions. For example, for every positive article written about a CEO in the *Wall Street Journal* or *Business Week*, the acquisition premium increased, on average, 4.8% (this translates to a large sum of money given that the acquisitions in this sample started at \$100 million). Hayward and Hambrick also showed that the premiums CEOs paid were not justified. The CEOs who received media praise, after all, might have been better CEOs than the managers they replaced; but on average, immediate returns to the acquisitions were  $-4\%$ , and 1-year returns were  $-11\%$ . Moreover, CEOs who paid higher premiums generally turned in worse performance as measured by immediate and 1-year returns.

Hayward and Hambrick (1997) treated their measures as indications of some underlying character trait, perhaps even pathology. A psychologist would likely explain the measures more as describing situations. Research suggests that most people are overconfident much of the time (e.g., Dunning et al., 1990; Fischhoff et al., 1977), but social situations often help to mitigate this potential hubris. Loved ones gently talk people out of crazy ideas, and peers give frank feedback when people overstep their abilities.

Thus, it is tragic but not surprising that CEOs, whose power means they have fewer social correctives for their own lack of self-knowledge, may be particularly prone to the perils of hubris.

They are most likely to make bad decisions when their social situations conspire to give them even more power—when a series of successes makes it appear that they are on a roll, when their next-closest colleagues have much less power (as illustrated by lower paychecks), or when the external media grant them more status by endorsing their past decisions.

In other research on acquisitions, Malmendier and Tate (2003) showed another effect that may be produced by a lack of self-insight. CEOs who held a large number of stock options in their own companies were more likely than other CEOs to acquire firms whose major lines of business differed significantly from their own. Economists classify mergers into two categories depending on the similarity of the firms that are being merged. *Related* mergers join two companies in the same general business (e.g., two electronics firms), whereas *diversifying* mergers join companies in different businesses (e.g., a railroad company and a steel producer). Compared with other CEOs, CEOs who are heavily invested in the stock of their own companies are more likely to conduct diversifying mergers—which indicates that they think their management skills extrapolate well to very different industries. The stock market, filled with independent observers who vote with their money, takes an especially dim view of such mergers—stock prices drop on average when any merger is announced, but diversifying mergers produce a reaction that is especially negative.

Interestingly, the problems of self-knowledge illustrated by overly optimistic CEOs may be apparent to other people (at least outside the CEOs' own organizations). Malmendier and Tate (2003) also found increased acquisitiveness by CEOs who were described by the financial press as “confident” and “optimistic,” as opposed to “reliable,” “cautious,” or “conservative.” These “confident” CEOs were 78% more likely to conduct a diversifying merger than less confident CEOs.

#### *Mitigating Overconfidence*

From our viewpoint, what is interesting about the problems of self-knowledge illustrated by Malmendier and Tate's (2003) CEOs is that the standard economic “solution” to the problem does not solve the problem and even may even make it worse. Traditionally, economists have worried that CEOs might engage in “empire building,” taking on too many projects or acquiring firms when they should not, because they are seeking to selfishly acquire additional status and prestige from managing a bigger firm. The standard solution to this problem is to give CEOs a bigger financial stake in their firms. Presumably, when CEOs own a significant amount of stock, they have a stronger incentive to make decisions that benefit all the firm's shareholders because they themselves are major shareholders.

But Malmendier and Tate (2003) found their largest effects on free cash flow, acquisitions, and diversifying acquisitions with CEOs who were heavily invested in their own firms. Our psychological analysis suggests that the problem is not one of incentives, but of self-knowledge. Presumably, Malmendier and

Tate's CEOs did not *intend* to waste their own money by undertaking diversifications, but thought they were doing something good for themselves and other shareholders. These results are consistent with psychological research on decision making, which indicates that adding incentives to a situation does not solve a problem that is caused by limited self-knowledge. Providing incentives for someone who is clueless simply ensures that he or she is both clueless and committed.

What, then, does mitigate CEO's problems with self-knowledge? One answer is, adding outside perspectives. It is not surprising that this works at the CEO level, given that many of the repairs that organizations develop for their employees involve forcing the employees to discuss decisions with other people who have an outside perspective (Heath et al., 1998). Hayward and Hambrick (1997) found that CEOs who scored high on their hubris measures were less likely to pay high acquisition premiums when there was an independent chairman of the board and when the board had a relatively high proportion of outsiders (i.e., directors who were not employees, suppliers, or buyers of the organization). Malmendier and Tate (2003) found that CEOs' tendency to acquire other firms is strongly mitigated by “effective corporate governance,” defined as having a board of directors that is not too large (in the corporate governance literature, this is considered to be under 12 members, a number that social psychologists might still consider too high for effective group discussions).

Another way of mitigating the problems is by taking away CEOs' access to money to invest. Because CEOs think the market is wrong about their prospects, they typically regard outside financing as too expensive and prefer to finance their projects internally. If the internal cash is taken away, CEOs behave more responsibly because they have to convince outsiders to loan them money for new projects. Organizations have many routines and procedures, such as loan limits, that repair front-line loan officers' tendency to be overly optimistic; a similar repair might be effective for CEOs.

It is important, however, that outsiders be able to maintain their outside perspective. Even “outside” advisors and consultants may be seduced into focusing on the inside view. Moore and Kim (2003) showed that when outsiders had the opportunity to bet on the outcomes of a particular individual, they made errors almost as large as that individual did. This finding suggests that outsiders may behave like insiders when they face common incentives. Interestingly, in Hayward and Hambrick's (1997) study, higher percentages of outside directors helped reduce acquisition premiums, but not the degree to which they held the firm's stock. This suggests that it is the outside perspective and not necessarily the incentives of the outsiders that improved judgment. Indeed, when officers and directors of the company held more stock, companies tended to pay higher premiums to make acquisitions, suggesting that when boards are highly motivated to increase the value of the firm, they may actually make worse decisions.