

Managing Overconfidence

J. Edward Russo • Paul J. H. Schoemaker

GOOD DECISION MAKING REQUIRES MORE THAN KNOWLEDGE OF FACTS, CONCEPTS, AND RELATIONSHIPS. IT ALSO REQUIRES METAKNOWLEDGE — an understanding of the limits of our knowledge. Unfortunately, we tend to have a deeply rooted overconfidence in our beliefs and judgments. Because metaknowledge is not recognized or rewarded in practice, nor instilled during formal education, overconfidence has remained a hidden flaw in managerial decision making. This paper examines the costs, causes, and remedies for overconfidence. It also acknowledges that, although overconfidence distorts decision making, it can serve a purpose during decision implementation. ☺

J. Edward Russo is Professor of Marketing and Behavioral Science at the Johnson Graduate School of Management, Cornell University. Paul J. H. Schoemaker is Associate Professor of Strategy at the Graduate School of Business, the University of Chicago. They recently published a book on managerial decision making, Decision Traps (Simon and Schuster, 1990.)

To know that we know what we know and that we do not know what we do not know, that is true knowledge.
—Confucius

Philosophers and writers have long tried to raise awareness about the difficulty of balancing confidence with realism, yet the consequences of unsupported confidence continue to plague businesses. Managers deal in opinions — they are bombarded with proposals, estimates, and predictions from people who sincerely believe them. But experience tells managers to suspect the certainty with which these beliefs are stated. For instance:

- A leading U.S. manufacturer, planning production capacity for a new factory, solicited a projected range of sales from its marketing staff. The range turned out to be much too narrow and, consequently, the factory could not adjust to unexpected demand.
- A loan officer at a major commercial bank felt that his colleagues did not understand their changing competition as well as they thought they did and were refusing to notice signs of coming trouble.
- In the early 1970s, Royal Dutch/Shell grew concerned that its young geologists too confidently predicted the presence of oil or gas, costing the company millions of dry-well dollars.

- The sales head for Index Technology, a new software venture, repeatedly received unrealistic sales predictions, not only on amounts but also on how soon contracts would be signed.

Managers know that some opinions they receive from colleagues and subordinates will be accurate and others inaccurate, even when they are all sincerely held and persuasively argued. Moreover, given any strongly held opinion, one seldom has to look far to find an opposing view that is held no less firmly. We do not even have to favor a position now to reserve the right to hold a *future* position. One of us attended a faculty meeting at which a senior faculty member had been notably silent during a heated debate. When asked for his position, he replied, "I feel strongly about this; I just haven't made my mind up which way."

People are often unjustifiably certain of their beliefs. As a case in point, the manufacturer cited above accepted the staff's confidently bracketed sales projections of twenty-three to thirty-five units per day and designed its highly automated factory to take advantage of that narrow range. Then, because of a worldwide recession, sales dropped well below twenty-three units per day. The plant was forced to operate far below its breakeven point and piled up enormous losses. Instead of being the best of the company's production facilities, it became the biggest loser.

Confidence Quiz

For each of the following questions, provide a low and a high estimate such that you are 90 percent certain the correct answer will fall within these limits. You should aim to have 90 percent hits and 10 percent misses. The correct answers are provided at the end of this article so that you can compute how close you come to the ideal level of one miss in ten.

90% Confidence Range

lower upper

1. How many patents did the U.S. Patent and Trademark Office issue in 1990?
2. How many of *Fortune's* 1990 "Global 500," the world's biggest industrial corporations (in sales), were Japanese?
3. How many passenger arrivals and departures were there at Chicago's O'Hare airport in 1989?
4. What was the total audited worldwide daily circulation of the *Wall Street Journal* during the first half of 1990?
5. How many master's degrees in business or management were conferred in the United States in 1987?
6. How many passenger deaths occurred worldwide in scheduled commercial airliner accidents in the 1980s?
7. What is the shortest navigable distance (in statute miles) between New York City and Istanbul?
8. What was General Motors' total worldwide factory sales of cars and trucks (in units) in the 1980s?
9. How many German automobiles were sold in Japan in 1989?
10. What was the total U.S. merchandise trade deficit with Japan (in billions) in the 1980s?

A Test of Confidence

How should managers deal with the often unreliable opinions they receive? The answer lies in recognizing that most people's beliefs are distorted by deep-seated overconfidence. Once we understand its nature and causes, we can better devise plans for controlling it. The first step is to document and measure the problem's severity.

Recall the loan officer who believed that his colleagues were overconfident about their competitors. He went to his boss with this concern and proposed measuring the degree of confidence his colleagues had in their knowledge about the bank's competitors. The boss insisted there was nothing to worry about: "No one is more realistic than a banker." Despite this overconfident answer, the boss agreed to the test — but only he would take it. To his surprise, he failed miserably; to his credit, he then asked all eleven other loan officers to take the same test. Every one of them flunked.

The loan officer's test asked for both best estimates and *ranges of confidence* around those estimates. The "Confidence Quiz" shown here is just such a test, one that involves general business rather than company-specific questions. The bank test included confidence

ranges for questions like: "What was the total dollar value of new commercial loans made by [Competitor X] last year?" and "What is the total number of commercial loan officers at the sixth- through tenth-ranking banks in [our city]?"

We believe that, like the loan officer's reluctant boss, you will be surprised at how poorly you do on the test below. For each of the ten quantities in the quiz, simply give a low guess and a high guess such that you are 90 percent sure the true value will lie between them. Try it before reading further, particularly if you are apt to be confident of your ability to predict accurately.

Metaknowledge

The confidence quiz measures something called *metaknowledge*: an appreciation of what we do know and what

we do not know. Normally, we define knowledge as consisting of all the facts, concepts, relationships, theories, and so on that we have accumulated over time. Metaknowledge concerns a higher level of expertise: understanding the nature, scope, and limits of our basic, or *primary knowledge*. Metaknowledge includes the uncertainty of our estimates and predictions, and the ambiguity inherent in our premises and world views.¹

At times, metaknowledge is more important than primary knowledge. For example, knowing when to see a lawyer or a doctor (metaknowledge) is more important than how much we know about law or medicine (primary knowledge). We draw on our metaknowledge when we conclude that we have enough information and are ready to make a decision now. If we think we are ready to decide when we are not, we may make costly mistakes. Only when we appreciate the limits of our primary knowledge can we sensibly ask for more or better information.

Examining confidence ranges, one of several ways researchers study metaknowledge, is a practical means of assessing personal uncertainty. Having sound metaknowledge means being able to predict within reasonable ranges.

Whether you should focus on 90 percent, 70 percent, or just 50 percent confidence ranges depends on

the issues and risks involved. When building a complex new oil refinery, where the downside risks are high, you may want to incorporate even extreme swings in oil prices. In that case, perhaps a 95 percent confidence range on future crude oil price levels should be assessed. However, for estimating regional sales levels, you may want a 50 percent range, as you can cope more easily with surprises outside that range.

We and others have found that whether managers are asked for 50 percent, 70 percent, or 90 percent confidence ranges, few employees are able to supply them realistically. Even experts, who by definition know a lot about a specialized topic, are often unable to express precisely how much they do not know. Yet to size up and factor uncertainty into our judgments is crucial to successful decision making. Experimental evidence suggests that this is a serious weakness in human judgment, even among those well versed in the use of quantitative tools. In technical language, few people are well calibrated; that is, few people can accurately assess their uncertainty. In business, this translates into risk underestimates, missed deadlines, and budget overruns.

Table 1 summarizes some results we have collected from different industries, most often using questions tailored to that industry and occasionally to a specific firm. No group of managers we tested ever exhibited adequate metaknowledge; every group believed it knew more than it did about its industry or company. Of the 2,000-plus individuals to whom we have given a ten-question quiz using 90 percent confidence intervals, fewer than 1 percent were not overconfident.

Our own evidence in Table 1 is confirmed by a large body of similar results from different professions, levels of expertise, and ages.² The only cross-cultural studies, done with Asian managers of several nationalities, further confirm the ubiquity of overconfidence.³

If a question falls outside your area of expertise, should you be excused if your confidence interval misses it? When we ask, "How many total employees did IBM have on its payroll on 31 December 1990?" managers outside the computer industry sometimes remark that IBM's staff size is irrelevant to their job, so they should be forgiven for their poor performance on an overconfidence quiz.

We strongly disagree. Whether you know a lot or a little about a subject, you are still responsible for knowing how much you don't know. If you know a lot, as a computer industry manager should, your 90 percent confidence intervals will be narrow; if you know less, they should be wider. In either case, your subjective 90 percent confidence intervals should, by definition, capture the true answers 90 percent of the time. (IBM had 365,000 employees on 31 December 1990.)

In actuality the job relevance of the questions does affect results, possibly because experience reduces overconfidence. In Table 1 we see that managers in the computer firm did better on firm-specific questions (58 percent misses) than on those covering their entire industry (80 percent misses). The data-processing managers, too, showed less overconfidence on industry-specific items (42 percent misses) than on general business facts (62 percent misses).

Although these results suggest that job relevance tends to reduce overconfidence, would such a pattern be confirmed by a more systematic study? Moreover, is the reduction in overconfidence only partial, or would questions very specific to people's jobs drastically reduce overconfidence? We tested this using ninety-six professionals drawn from a variety of occupations. We used two different confidence quizzes: the first contained fifteen job-specific questions; the second contained fifteen questions unrelated to these professionals' jobs. The unrelated questions were created by having pairs of professionals exchange questionnaires. Thus, one person's job-specific questions become the other's unrelated questions and vice versa. As a check, we asked everyone afterwards to rate the job relevance of the fifteen questions in the

Table 1 Overconfidence across Industries

Industry Tested	Kind of Questions Used in Test	Percentage of Misses		Size**
		Ideal*	Actual	
Advertising	Industry	10%	61%	750
	Industry	50	78	750
Computers	Industry	5	80	1290
	Firm	5	58	1290
Data processing	Industry	10	42	252
	General business	10	62	261
Money management	Industry	10	50	480
Petroleum	Industry & firm	10	50	850
	Industry & firm	50	79	850
Pharmaceutical	Firm	10	49	390
Security analysis	Industry	10	64	497

* The ideal percentage of misses is 100% minus the size of the confidence interval. Thus, a 10% ideal means that managers were asked for 90% confidence intervals.

** The total number of judgments made across persons and questions.

first quiz, using a scale from 1 (irrelevant) to 7 (highly relevant). For these 90 percent confidence ranges, the unrelated quiz yielded 53 percent misses (instead of the ideal 10 percent). For the job-relevant quiz, the percentage of misses went from a high of 58 percent for the least relevant questions to 39 percent for the most relevant ones. Figure 1 displays this downward trend.⁴ Note that overconfidence does not vanish, but remains at 29 percent over the ideal, even for the most relevant questions.

In sum, better primary knowledge is generally associated with better (though still imperfect) metaknowledge. That is, experts know better what they don't know, and this fact is one key to effective solutions, as we discuss next.

Developing Good Metaknowledge

How might professionals develop a sharper sense of how much they do and do not know? Once the existence of overconfidence is acknowledged, two elements are essential: *feedback* and *accountability*.

Feedback that is accurate, timely, and precise tells us by how much our estimates missed the mark. Accountability forces us to confront that feedback, recalibrate our perceptions about primary knowledge, and temper our opinions accordingly.

One mistake we often see managers make is equating experience and learning. *Experience is inevitable; learning is not.* Overconfidence persists in spite of experience be-

cause we often fail to learn from experience.⁵ In order to learn, we need feedback about the accuracy of our opinions and doubts. We also need the motivation to translate this information into better metaknowledge.

At least three groups of professionals have used systematic feedback and accountability to develop excellent metaknowledge: Shell's geologists, public accountants, and weather forecasters.

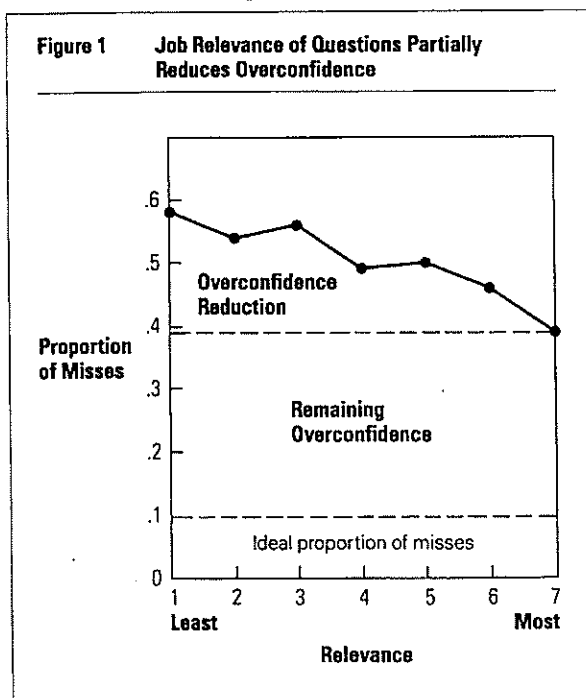
- **Shell's Geologists.** Recall the earlier example of Royal Dutch/Shell, the Anglo-Dutch oil and gas giant. Shell had noticed that newly hired geologists were wrong much more often than their levels of confidence implied. For instance, they would estimate a 40 percent chance of finding oil, but when ten such wells were actually drilled, only one or two would produce. This overconfidence cost Shell considerable time and money.

These judgment flaws puzzled senior Shell executives, as the geologists possessed impeccable credentials. How could well-trained individuals be overconfident so much of the time? Put simply, their primary knowledge was much more advanced than their metaknowledge. To develop good metaknowledge requires repeated feedback, which was coming too slowly and costing too much money.

In response, Shell designed a training program to help geologists develop *calibration power*. As part of this training, the geologists received numerous past cases that incorporated the many factors affecting oil deposits. For each case, they had to provide best guesses as well as ranges that were numerically precise. Then they were given feedback as to what had actually happened. The training worked wonderfully: now, when Shell geologists predict a 40 percent chance of producing oil, four out of ten times the company averages a hit.

- **Public Accountants.** When experienced auditors provided estimates and confidence ranges for account balances, they actually proved slightly underconfident.⁶ Their ranges were too wide rather than too narrow. Perhaps accountants have learned to compensate for overconfidence because of their role as detectors of fraud and error. The profession places an extraordinarily high value on conservative judgments.

- **Weather Forecasters.** But what, then, are we to make of the only other professional group that has been found to be well calibrated: U.S. Weather Service forecasters? Figure 2 tells a remarkable story of accurate subjective probabilities.⁷ When U.S. Weather Service forecasters predicted a 30 percent chance of rain, as they did 15,536 times in this study, it rained almost exactly 30 percent of the time. This superb accuracy holds along the entire range of probability, except at the highest levels. When a

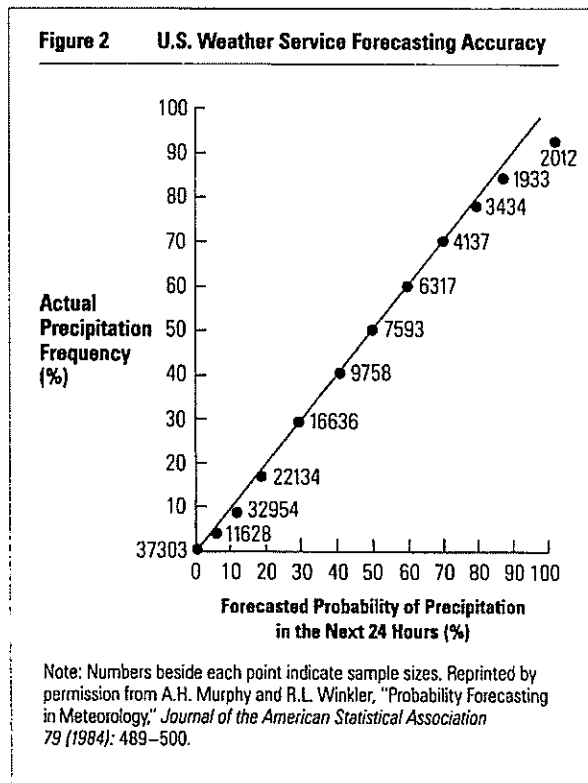


100 percent chance of rain' is predicted, it actually rains only 90 percent of the time. This prediction error reflects deliberate caution on the part of the forecasters.

What these three groups have in common is precise, timely feedback about repeated judgments in a field whose knowledge base is relatively stable, unlike the stock market or fashion industry, for instance. Furthermore, all three groups are held accountable by their supervisors or professional colleagues for the accuracy of their confidence judgments. Within a day, the weather forecasters receive feedback about whether or not it rains, and their predictive performance is factored into their salary increases and promotions. We believe that timely feedback and accountability can gradually reduce the bias toward overconfidence in almost all professions. *Being "well calibrated" is a teachable, learnable skill.*

Organizations can accelerate the slow, costly process of learning from experience by keeping better track of managerial judgments and estimates. Performance reviews should emphasize the value to the firm of realism and back this emphasis up with both assessments and incentives. In addition, training programs can provide feedback on simulated or past decisions whose outcomes are not widely known, just as Shell's training program did.

Systematic feedback works, even though it treats only



the symptoms of overconfidence. That is, it corrects overconfidence without teaching what caused it in the first place. Several other techniques for reducing overconfidence directly attack its causes. No single cause or prototypical situation can be consistently connected with overconfidence. There are three classes of causes: *cognitive, physiological, and motivational.*

Cognitive Causes of Overconfidence

- **Availability.** A major reason for overconfidence in predictions is that people have difficulty in imagining all the ways that events can unfold. Psychologists call this the availability bias: what's out of sight is often out of mind.⁸ Because we fail to envision important pathways in the complex net of future events, we become unduly confident about predictions based on the fewer pathways we actually do consider.

The limited paths that are evident (e.g., the expected and the ideal scenarios) may exert more weight on likelihood judgments than they should. Bridge players provide a telling example of how availability can cause overconfidence.⁹ More experienced bridge players are better-calibrated bidders because they take into account more unusual events or hands. Less experienced players believe they can make hands they often cannot, precisely because they fail to consider uncommon occurrences.

- **Anchoring.** A second reason for overconfidence relates to the anchoring bias, a tendency to anchor on one value or idea and not adjust away from it sufficiently.¹⁰ It is typical to provide a best guess before we give a ballpark range or confidence interval. For example, we usually estimate next quarter's unit sales before we come up with a confidence range. The sales estimate becomes an anchor point and drags the high and low brackets, preventing them from moving far enough from the best estimates.

How strong is the best estimate's pull? To answer this question, we presented twenty trivia questions (e.g., "What is the length of the Nile River?") to two groups of managers. One group (of eighty-four people) first gave a best estimate, that is, an anchor point, and then provided a 90 percent confidence interval. The second group (of fifty-one people) directly supplied a confidence range without ever committing to a best guess. The first group scored 61 percent misses (compared to the ideal of 10 percent). In contrast, the unanchored group's intervals were wider and missed only 48 percent of the true answers. Thus, overconfidence was reduced substantially by simply skipping best guesses and moving directly to ranges. (The Nile River is 3,405 miles long.)

Although this de-anchoring technique has yet to be

verified outside the controlled laboratory, we see no reason why it should not work as well in managerial environments. Interestingly, how well it works will depend on managers' ability to focus on the confidence interval and block out of their thinking any earlier estimate that might serve as an anchor.

- **The Confirmation Bias.** A third cognitive reason for overconfidence concerns our mental search process. When making predictions or forecasts, we often lean toward one perspective, and the natural tendency is to seek support for our initial view rather than to look for disconfirming evidence. Unfortunately, the more complex and uncertain a decision is, the easier it is to find one-sided support. Realistic confidence requires seeking disconfirming, as well as confirming, evidence.¹¹

How much weight to give to evidence, pro or con, is a complex issue depending both on the strength of the evidence itself and on the credibility of its source. Griffin and Tversky, for instance, suggest that people over-weight the strength of evidence (e.g., how well a candidate did in an interview) relative to the credibility of that type of evidence (the limited insight gained from any single interview).¹² Whenever source credibility is low, as is often the case in business, and the strength of the evidence is highly suggestive, overconfidence is likely to occur. Thus, the interviewed candidate is too confidently predicted to be a winner or loser, given the fallible, limited evidence obtainable from a short, one-time interview. Ironically, Griffin and Tversky predict underconfidence under reverse circumstances, when the credibility of the source is high, but the evidence does not point strongly to one action or opinion.

- **Hindsight.** Hindsight makes us believe that the world is more predictable than it really is. What happened often seems more likely afterwards than it did beforehand, since we fail to appreciate the full uncertainty that existed at the time. Recall, for instance, George Bush's landslide victory over Michael Dukakis in 1988 (54 percent of the popular vote). At the time of the nominating conventions, the outcome of the election seemed far from certain. Indeed, University of Chicago MBA students that summer gave Bush only a 49 percent chance of winning,¹³ and the political press frequently cited the "wimp" factor. Nonetheless, the results might seem quite predictable some years later. Hindsight instills an illusion of omniscience.

Cognitive Remedies to Overconfidence

What remedies are available for cognitive sources of overconfidence? We look at five techniques.

- **Accelerated Feedback.** Recall how successful Shell was in training its junior geologists on past cases where the outcome was known, so they could get immediate feedback. Simple experiments with trivia questions have demonstrated the efficacy of accelerated feedback.¹⁴ This kind of training can be especially effective for new employees. Using tests derived from actual company records, employees could be trained to estimate their confidence in knowledge relevant to their new jobs. At first, these predictions will almost certainly be overconfident, but good feedback will quickly reduce it. And, in contrast to learning from experience, which tends to be slow and expensive, good feedback will reduce overconfidence cheaply.

But what can you do when faced with a single decision that you must make soon? Try to improve your thinking by bringing to mind relevant considerations that might easily be overlooked. The next four techniques offer specific methods for doing so.

- **Counterargumentation.** Think of reasons why your initial beliefs might be wrong, or ask others to offer counterarguments. Several studies have demonstrated the power of generating counterarguments, including one where a major company tested its managers by asking questions such as the following:¹⁵

Our company's current liabilities (defined as notes payable, short-term loans, etc.) were \$1,971 million and \$1,551 million as of December 31, 19xx, and March 31, 19xx, respectively. For October 31, 19xx, the company's current liabilities will be (circle one):

(a) *greater than \$1,900 million*

(b) *less than or equal to \$1,900 million*

Give your subjective probability that you will be correct: _____ %.

Half of the participants in the experiment were merely asked to circle (a) or (b) and then state how confident they were about their choice. This group's mean estimated probability of being correct was 72 percent. However, they actually picked the correct answer only 54 percent of the time. Hence, they were overconfident by 18 percent. The other half of the participants were asked to think of "the major reason why the alternative circled might be wrong" before giving their subjective probabilities. That is, they were asked to think disconfirmingly and provide at least one counterargument to their initial guess. Then they were allowed to change their answer if they wished. This group's average estimated probability of being

correct was 73 percent, and they actually picked the correct answer 62 percent of the time. Thus, their level of overconfidence was only 11 percent, *a reduction of nearly two-fifths thanks to a single counterargument*. Other studies have found that, when listing pros and cons, the cons do the most good in countering overconfidence.¹⁶

But is this practical? It depends. We see no reason why major capital budgeting requests could not have a counterargument section in which managers are asked to identify the major reasons not to go ahead. And if the project does fail later, the actual causes had better be listed in this contrarian section of the report. A warning, however. To be useful, this process must be taken seriously, with serious consequences (both good and bad) for the managers involved. Otherwise it may degenerate into a useless formality or, worse, corporate "gaming" (i.e., managers withholding their genuine concerns in favor of saying whatever gets the budget approved). A second problem is that managers may truly not recognize potential hazards. The next tactic, Paths to Trouble, addresses that problem.

• **Paths to Trouble.** If we are overconfident in predicting success because we cannot see the paths to potential trouble, fault trees may help. A fault tree is a hierarchical diagram designed to help identify all the paths to some specific "fault" or problem. For an example, see Figure 3. To be useful, fault trees must be reasonably complete, at least in identifying the major categories of potential trouble. If they are not, chances are that even specialists will fail to realize what is missing.¹⁷ People assume the causes listed account for almost everything that could go wrong and underestimate the final category, "all other" causes of failure. In the study that used the restaurant fault tree, the "all other" category was estimated by hospitality industry managers to contain only 7 percent of the chance that something might go wrong, whereas it really contained 54 percent of the chance.¹⁸

How can this blindness be overcome? Warning people does not seem to help.¹⁹ What does work is asking people to extend the fault tree by listing additional causes of the problem.²⁰ In the restaurant study, some people were shown branches with only six instead of twelve causes and asked to provide further reasons. When two more causes were listed by managers themselves, the original omission error of 46 percent dropped to 23 percent; when four were listed, it dropped to 12 percent; and when six possible causes were added, it disappeared entirely. In sum, *the more causes generated, the smaller is the error of assuming that all relevant causes are already listed*.

• **Paths to the Future.** If deeper thinking is called for, beyond the listing of reasons, explicit scenario analysis may be useful. Whereas fault trees highlight individual causes, scenarios focus on their conjunction. Scenarios are script-like narratives that paint in vivid detail how the future might unfold in one or another direction. Envisioning vastly different worlds than those expected has helped companies like Royal Dutch/Shell to better estimate economic and political uncertainty.²¹ A direct test that compared 90 percent confidence intervals before and after scenario construction found, on average, a 30 percent stretching of ranges.²² Asking managers to construct different scenarios makes them better appreciate the uncertainty in key parameters or estimates. In addition, it often provides new ideas for innovation or competitive positioning.

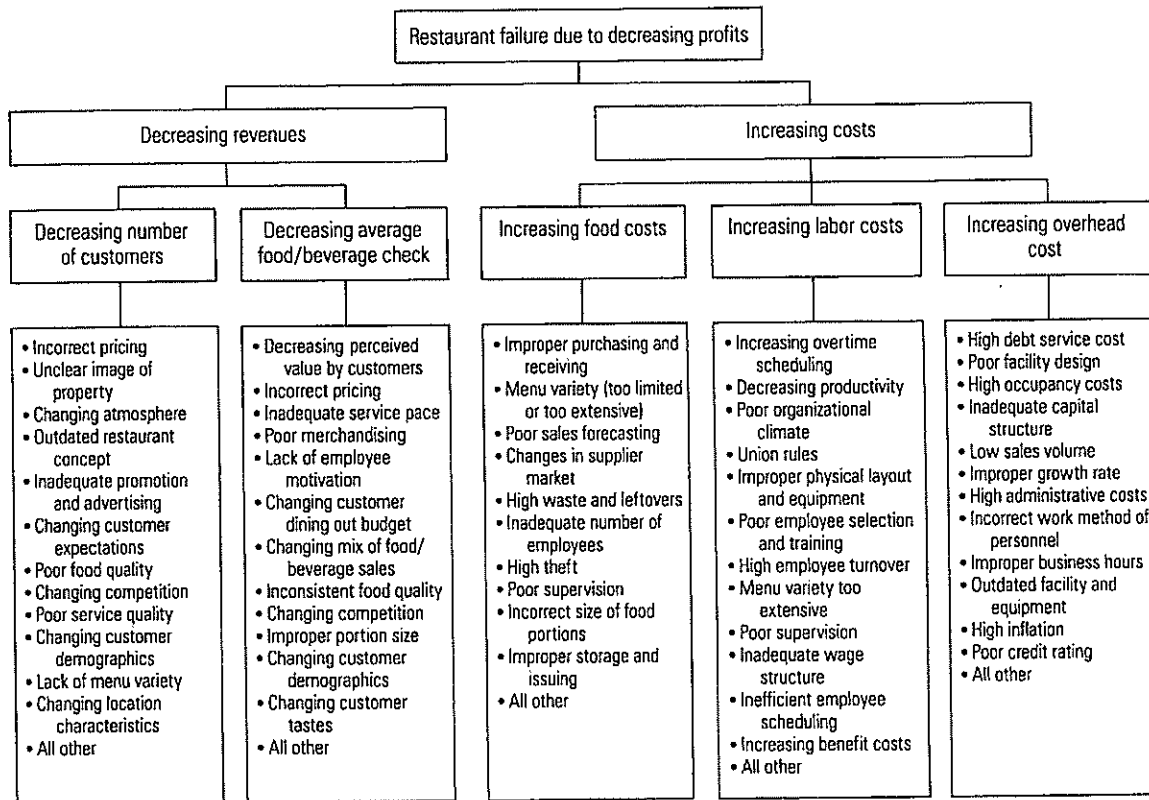
• **Awareness Alone.** Although these techniques are valuable, we happily acknowledge that, for many managers, awareness alone may be all that is needed. Good managers often devise their own solutions to the problems of overconfidence.

Recall the head loan officer who was certain that he and his staff knew their competitors quite well. After failing the tailored overconfidence quiz, he took immediate action. Each officer was required to contribute information to a "competitor alert" file. And each was required to check the file weekly to gain a more realistic appreciation of their competition. Within three weeks, a loan officer found information in the file signaling that a major client was contemplating a shift to another bank. The competitor was not one of the city's major commercial banks, and the magnitude of the business at risk exceeded the competitor's legal lending limit. However, by joining with another institution, it was able to offer a loan large enough to meet the client's needs. Thus alerted, the loan officer in charge of the account convinced the client not to switch banks, saving \$160,000 in annual revenue.

The head of sales for Index Technology took a different approach when his salespeople were overconfident about if and when orders of the company's product would be written by potential customers. He called some customers himself. His salespeople didn't like it, but the approach worked: soon they were predicting orders and the timing of those orders much more accurately.

A negotiation experiment further underscores the value of awareness alone.²³ Subjects believed they had a 65 percent chance of winning in a simulated negotiation task entailing binding arbitration. In this set-up, for every winner there had to be a loser, implying a 50

Figure 3 Fault Tree for Restaurant Failure



Reprinted by permission of John Wiley & Sons, Ltd., copyright 1988. From L. Dubé-Rioux and J.E. Russo, "An Availability Bias in Professional Judgement," *Journal of Behavioral Decision Making* 1 (1988): 223-237.

percent chance of winning. Hence, most of the people were overconfident. The researchers then took a random half of the people aside and warned them about overconfidence. Compared to the unwarned group, those forewarned were 30 percent more likely to reach a negotiated agreement instead of having to turn to costly arbitration, and they achieved net dollar benefits that were 70 percent higher. As skilled lawyers know, negotiating a settlement is one area where realism pays.

General Versus Specific Awareness

Although general awareness of a bias is invaluable, it does not guarantee that the bias will be spotted in every instance. Consider this study. Twelve financial officers were asked to estimate ten quantities pertinent to their

organization's business operations and to provide a 90 percent confidence interval for each.²⁴ As usual, these intervals failed to capture the true value a high percentage of the time; in this case, the failure rate was 78 percent versus the ideal of 10 percent. In addition, the financial officers were asked to estimate "how many of the ten intervals you gave . . . will contain the actual value." This was asked immediately after the ten intervals were constructed and before the true answers were revealed. Of course, every officer should have answered "nine" because that is what a 90 percent confidence interval means, by definition. However, only one did. The others estimated that fewer than nine of the ten intervals would capture the true value. On average, the twelve officers guessed that they would miss 5.6 of the ten questions but couldn't tell which ones. These data suggest that people are more aware of overconfidence in general

than they are in particular.

The same problem — of general awareness but specific blindness — was described by John Stuart Mill, the 19th century economist and social philosopher, in *On Liberty*:

Unfortunately for . . . mankind, the fact of their fallibility is far from carrying the weight of their practical judgment, which is always allowed to it in theory; for while everyone knows himself to be fallible, few . . . admit the supposition that any opinion of which they feel very certain may be one of the examples of the error to which they acknowledge themselves to be liable.

Physiological Causes of Overconfidence

Because overconfidence is a distortion of judgment, it is often thought of as a purely mental phenomenon; however, at times it has biochemical causes. Euphoria, the elated feeling of well-being that commonly follows personal or professional success, may cause overconfidence. (The biochemical compounds involved in euphoria appear to be hormones, such as adrenalin and endorphines, that the body produces as a response to strong emotional reactions.) We also suspect that drugs like cocaine and alcohol can produce overconfidence.²⁵

Ford Motor Company provides an example of how a major firm dealt with the negative side effects of euphoria. As the 1970s ended, Ford faced hard times: reduced market share, layoffs, and the superior quality of Japanese cars. In response to these conditions, Ford organized meetings of its plant managers and assistant managers to solicit and communicate suggestions for improving manufacturing quality. The flood of ideas, and the resulting enthusiasm for what might be accomplished, swept nearly everyone away. Wisely, Ford's top management imposed a cooling-off period of several weeks before the returning managers could implement any of the suggestions. Because of the euphoric mood at the end of the meetings, senior executives distrusted their managers' judgment, and they wanted time for a more calculated look prior to committing major funds.

Dealing successfully with physiological causes of overconfidence, as with all types of overconfidence, requires awareness of the problem: you can't fix it if you can't find it. In this regard, individual awareness is the single most important factor. If you are euphoric, wait

to commit yourself to a plan of action, just as, if you drink, don't drive.

Overconfidence in Group Judgments

By this time, you may wonder if groups do better than individuals when sizing up uncertainty. The answer is mixed.²⁶ Group judgments can be better than individual ones, precisely because in groups people are forced to recognize that others see the world differently than they do. This often sparks a realization that perhaps their own views are held with unjustifiable conviction. At other times, however, groups may bolster the majority opinion to even more extreme levels.

To test group overconfidence, we conducted a simple experiment with eighty-three managers. First, people were asked to privately form 90 percent confidence ranges on ten questions. Then they were asked to compare and discuss their ranges in groups of three or four in order to come up with a single group range for each question. We did not specify how this was to be accomplished. Some groups argued heatedly; others merely averaged the individual guesses; still others used the most extreme values in the group as their outer brackets. After the group decisions were made, people were allowed to change their private ranges. The initial, unrevised private judgments generated an average of 72 percent misses (compared to an ideal of 10 percent), signifying serious overconfidence. Group judgments were significantly less overconfident, with 56 percent misses on average. The revised private judgments made after the group decisions resulted in 62 percent misses.

Making Group Judgments Better

On average, the group judgments were better than individual judgments in the above task. At worst, they forced a compromise; at best, they encouraged openmindedness. Individually, however, people may still anchor too strongly to their initial view and return to it when given the chance. This stubbornness can be to their, and their company's, detriment. There are relatively simple techniques for minimizing this recidivist tendency.

Delphi techniques and other procedures for sharing and averaging opinions are especially feasible in a networked PC environment. Rather than go around the table, collect people's initial estimates and ranges privately. Next, share these ranges and only then commence a debate. After group discussion, ask for one

more round of opinions and run with those averages. An extensive literature exists on expert aggregation and averaging, as well as experimental work on individual versus group calibration.²⁷

Motivational Factors in Overconfidence

Overconfidence isn't all bad! One legitimate cause of overconfidence is our need to believe in our abilities. Indeed, confidence in one's abilities is particularly widespread. At the beginning of a course, we often ask our MBA students (anonymously) whether their final grade will be in the bottom or top half of the class. The great majority are certain they will finish in the top half, *and they are willing to bet on it.*

Many of these people are distorting reality, yet their optimism has motivational value. Would risky projects be undertaken if a few key people did not have an unrealistic belief in their chances of success? As Goethe wrote, "For a man to achieve all that is demanded of him he must regard himself as greater than he is."

If the motivating value of overconfidence is clear, so is its downside. The value and danger of overconfidence may especially conflict for entrepreneurs. They often take risks others would not, and they must persuade investors and employees to join them in highly uncertain endeavors. Yet their eventual success also requires realism. A partner at a venture capital firm summarized this problem: "You expect entrepreneurs to have . . . an unshakable sense that they absolutely cannot fail. Yet since we will be partners with these people, we want to be sure that their egos will not stand in the way of making the best decisions for the business."

Moreover, to succeed in many business endeavors, we have to project confidence even when it cannot be justified. Because people often equate confidence with competence, you had better sound confident if you want your opinions to be treated as credible. It is seldom easy to stand up at an important meeting and say, "I'm not sure." Instead, people go out on a limb.

Can anything be done to reconcile the danger of distorting reality with the value of optimism? Perhaps the best advice is: Don't fool yourself. Don't permit yourself to be overconfident when making important decisions or commitments.

Deciding and Doing

We believe that much of the damage can be avoided if managers distinguish between deciding and doing. Deciding requires realism. But in implementing the

decision, the motivational benefits of overconfidence frequently outweigh its dangers.

Separating deciding from doing is not simple, and it's harder than it used to be. A century ago, business organizations were more like vertical pyramids: the deciders were on top and the workers were underneath. But in today's "flattened" organizations, every manager is both a decider and a doer.

So what should today's managers do? Our recommendation is to *be aware of when you are functioning as a decider and when your primary role is that of doer, motivator, or implementer.* When you are deciding, be realistic, both about how much you know and how much you don't know. When you are implementing, indulge in overconfidence when, and if, it is valuable to your performance or that of others.

All of us need self-confidence to function. We might not show up for work every morning if we did not believe we could make a difference. Nevertheless, too much confidence can backfire — can cause us to bet on plans, people, or projects which a more realistic appraisal would have rejected. Though normally an advocate of rational calculation, Lord Keynes keenly observed this human dilemma: "A large proportion of our positive activities depend on spontaneous optimism rather than on mathematical expectation . . . if animal spirits are dimmed and the spontaneous optimism falters, leaving us to depend on nothing but mathematical expectation, enterprise will fade and die." ♦

References

The authors acknowledge Janet Sniezek and Ilan Yaniv for their constructive comments and Jack B. Williams for his editorial advice.

1. Linguists distinguish between language competence (the ability to produce coherent statements) and metalinguage (the ability to state the rules of the language). Such a clear distinction does not always exist between primary knowledge and metaknowledge. Early in the century, U.S. Weather Service forecasters simply predicted whether or not it would rain (a statement of their primary knowledge). Now they provide an explicit probability of rain, making uncertainty assessment an explicit part of their primary knowledge.
2. S. Lichtenstein, B. Fischhoff, and L.D. Phillips, "Calibration of Probabilities: The State of the Art to 1980," in *Judgment under Uncertainty: Heuristics and Biases*, eds. D. Kahneman, P. Slovic, and A. Tversky (New York: Cambridge University Press, 1982), pp. 306-334.
3. G.N. Wright and L.D. Phillips, "Cultural Variations in Probabilistic Thinking: Alternative Ways of Dealing with Uncertainty," *International Journal of Psychology* 15 (1980): 239-257.
4. All claims made about differences or trends are statistically significant at the .05 level or lower. The sample sizes for the percentages in Figure 1 range from a low of 122 when relevance = 1 to a high of 270 when relevance = 7, with the unrelated percentage based on all 1,440 unrelated questions.

5. J.E. Russo and P.J.H. Schoemaker, *Decision Traps* (New York: Simon and Schuster, 1990).
6. L.A. Tomassini et al., "Calibration of Auditors' Probabilistic Judgments: Some Empirical Evidence," *Organizational Behavior and Human Performance* 30 (1982): 391-406.
7. A.H. Murphy and R.L. Winkler, "Probability Forecasting in Meteorology," *Journal of the American Statistical Association* 79 (1984): 489-500.
8. A. Tversky and D. Kahneman, "Availability: A Heuristic for Judging Frequency and Probability," *Cognitive Psychology* 4 (1973): 207-232;
- B. Fischhoff, P. Slovic, and S. Lichtenstein, "Fault Trees: Sensitivity of Estimated Failure Probabilities to Problem Representation," *Journal of Experimental Psychology: Human Perception and Performance* 4 (1978): 330-344.
9. G. Keren, "Facing Uncertainty in the Game of Bridge: A Calibration Study," *Organizational Behavior and Human Decision Processes* 39 (1987): 98-114.
10. P. Slovic and S. Lichtenstein, "Comparison of Bayesian and Regression Approaches to the Study of Information Processing in Judgment," *Organizational Behavior and Human Performance* 6 (1971): 641-744;
- A. Tversky and D. Kahneman, "Judgment under Uncertainty: Heuristics and Biases," *Science* 185 (1974): 1124-1131.
11. J. Klayman and Y.W. Ha, "Confirmation, Disconfirmation, and Information in Hypothesis Testing," *Psychological Review* 94, 2 (1987): 211-228.
12. D. Griffin and A. Tversky, "The Weighing of Evidence and the Determinants of Confidence" (Waterloo, Ontario: University of Waterloo, working paper, 1991).
13. P.J.H. Schoemaker, "Scenario Thinking" (Chicago: Graduate School of Business, University of Chicago, working paper, 1991).
14. For a review, see Lichtenstein, Fischhoff, and Phillips (1982).
15. J. Mahajan and J.C. Whitney, Jr., "Confidence Assessment and the Calibration of Probabilistic Judgments in Strategic Decision Making" (Tucson: University of Arizona, working paper series #12, 1987).
16. S.J. Hoch, "Availability and Inference in Predictive Judgment," *Journal of Experimental Psychology: Learning, Memory, and Cognition* 10 (1984): 649-662.
17. Fischhoff, Slovic, and Lichtenstein (1978).
18. L. Dubé-Rioux and J.E. Russo, "An Availability Bias in Professional Judgment," *Journal of Behavioral Decision Making* 1 (1988): 223-237. In this study, six of the twelve listed causes in a branch of a fault tree (see Figure 3) were removed. If people, in this case hospitality industry managers, were properly aware of all the major causes, then all of the probability of these six unlisted causes should show up in the last, "all other" category. In fact, very little did, strongly suggesting that what is out of sight is out of mind; i.e., the availability bias operates.
19. Fischhoff, Slovic, and Lichtenstein (1978).
20. Dubé-Rioux and Russo (1988).
21. P. Wack, "Scenarios: Uncharted Waters Ahead," *Harvard Business Review*, September-October 1985, pp. 73-89;
- P. Wack, "Scenarios: Shooting the Rapids," *Harvard Business Review*, November-December 1985, pp. 139-150.
22. Schoemaker (1991).
23. M.A. Neale and M.H. Bazerman, "The Effects of Framing and Negotiator Overconfidence on Bargaining Behavior and Outcomes," *Academy of Management Journal* 28 (1985): 34-49.
24. J.A. Sniezek and T. Buckley, "Level of Confidence Depends on Level of Aggregation," *Journal of Behavioral Decision Making* 4 (1991): 263-272.
25. We wonder how many traffic fatalities are caused by alcohol-induced overconfidence. Certainly driving skills are impaired by alcohol, but this may be only part of the story. A more deadly aspect is that the drinker's confidence is not reduced nearly as much as the ability itself. This confidence gap between the skill levels drivers believe they possess and the reduced levels they actually have seems to be a primary problem with drunk drivers.
26. Despite a presumption that "two heads are better than one," groups do not always make better decisions than individuals. The phenomenon known as groupthink is one serious problem. Whether groups are superior seems to depend on whether conflict is articulated or swept under the rug. See:
 1. I.L. Janis, *Groupthink*, 2nd ed. (Boston: Houghton Mifflin, 1982.)
 2. R.T. Clemen and R.L. Winkler, "Unanimity and Compromise among Probability Forecasters," *Management Science* 36 (1990): 767-779; and
 3. J.A. Sniezek and R.A. Henry, "Accuracy and Confidence in Group Judgment," *Organizational Behavior and Human Decision Processes* 43 (1991): 1-28.

The answers to the Confidence Quiz: (1) 96,727 patents; (2) 111 Japanese corporations; (3) 59,130,007 arrivals and departures; (4) 2,076,713; (5) 67,496 degrees; (6) 6,700 deaths; (7) 5,757 miles; (8) 77.8 million units; (9) 147,324 automobiles; and (10) \$354 billion.

Reprint 3321