J. Edward Russo, Paul J.H. Schoemaker. Decision Traps:
Ten barriers to brilliant decision-making and how to overcome them.
Doubleday 1989 ISBN: 0385238350 (280 total pages)

Chapter 5
Improving Your Intelligence-Gathering

> ***It's not what we don't know that gives us trouble, it's what we know that ain't so.***
>
> WILL ROGERS

How can you obtain really good intelligence when you and your staff, like all humans, suffer from serious biases in your thinking?

Mere awareness of the problem helps. An educated manager should seek to avoid overconfidence about judgments, limit availability biases by gathering information systematically, and minimize the drag of anchoring.

But the tendency to these errors so pervades human thinking that you may need more than mere awareness to overcome them. This chapter introduces five fairly simple methods that can help transform your intelligence-gathering from hit-or-miss into a reliable part of an excellent decision-making process.

> ***Decision Trap Number 5:***
>
> **Shortsighted Shortcuts**—Relying inappropriately on "rules of thumb" such as implicitly trusting the most readily available information or anchoring too much on convenient facts.

Step One: Require Every Estimate to Include a "Level of Confidence"

The first step in making intelligence-gathering really systematic is to ask: "What does an estimate mean?" Until you know exactly what estimates mean, you can do little to improve them.

Suppose your sales manager announces that he is "reasonably sure" he can sell 200,000 Whatzits next year. You know that doesn't mean he will sell *exactly* 200,000. Should the manufacturing department plan its production for *exactly* 200,000? Of course not. Then what does his estimate tell you? Indeed, does the sales manager himself know exactly what he means?

To create really useful estimates, you need to put them in a form whose meaning can be clearly understood—a form that allows you to evaluate whether a forecaster's record is good or bad. The best way is to state a range and a level of confidence: "I estimate we will sell between 180,000 and 220,000 Whatzits next year. I am 80 percent sure." This kind of estimate can be far more useful than a numerically exact (but not exactly correct) prediction. It allows your organization to plan the future with as much knowledge as currently exists or invest in getting better information in case the uncertainty is too great. And it enables the forecaster to learn much more easily from experience.

The sales manager who forecasts a range and a confidence level is betting his reputation in a way that he is not when he simply says he is "reasonably sure" he will sell 200,000 Whatzits. Suppose that for three years in a row he says he is reasonably sure he will sell 200,000. In reality, sales are 175,000 the first year, 140,000 the second, and 230,000 the

third. If he had just given a single-point forecast, your staff could argue endlessly over whether his mistakes were reasonable.

But had he forecast a range of 180,000 to 220,000 each year with a confidence level of 80 percent, you can be pretty sure that something is wrong: Since three years in a row sales failed to fall in the range that was forecast, the sales manager should feel embarrassed. If you are his boss, you'll know that either he needs training in better forecasting or the company should focus on shorter-term forecasts, which should be easier for him.

Of course, erroneous forecasts not only hurt your sales manager's reputation, they may hurt your company as a whole. To protect his reputation, your sales manager will probably learn to forecast confidence ranges more accurately. After a while, you may find that when he says he is "80 percent sure" of something, there's really an 80 percent chance it will happen. Ideally, the sales manager will produce his estimates using the linear model approach we'll discuss in Chapter 6.

Now the sales manager's estimates will enable the manufacturing department to deal realistically with suppliers and the finance department to project realistic cash flows. Without confidence-range estimates, however, good risk assessments are not really possible.

Fine-Tuning Confidence Judgments

Sometimes you need not state a prediction in terms of a range. If, for example, you are recommending the selection of a contractor and you just want to indicate that you are confident the person will show up on a particular day, you may simply say you are 95 percent sure that the firm will show up and do an acceptable job.

But even in this kind of case, providing a confidence level helps make judgments useful. Both the people making the forecast and the person relying on it know that forecasters can't say they're 95 percent sure ten times and be wrong eight of the ten without facing consequences.

And providing levels of confidence will help you analyze whether you need to seek *additional* information before making a decision on a complex issue. When a confidence-range estimate has been made you can ask, "What additional information would allow us to predict a narrower range with an even higher level of confidence?" Narrower ranges and higher levels of confidence mean less uncertainty for your business. Thus, you will be better able to assess whether the additional information and delay are worth the added costs.

Watch out in employing this tactic with analysts who are not used to it: The analysts may *feel* more confident with more information even if the additional information does not make them any more *accurate* at all. (We'll discuss this point in greater detail later in this chapter.) But as people get used to using confidence-range forecasts, they are more likely to be able to tell you accurately what would help them forecast better—or to say honestly that they can't be certain.

Step Two: Provide Feedback and Training to Help People Calibrate

Once you've introduced ranges and confidence levels in estimates, you need to ensure that the people who make estimates receive timely feedback on their accuracy, that they are trained to improve, and that they are accountable for excellence.

Royal Dutch/Shell, for instance, has saved many millions of dollars by improving the accuracy of oil-drilling decisions through training in prediction skills. After Shell began requiring that geologists state their levels of confidence in their

predictions, it discovered (not surprisingly) that newly hired geologists were wrong much more often than their stated levels of confidence predicted. When young geologists said there was a 50 percent chance of finding oil, oil would be found as little as 30 percent of the time. Thus, Shell was spending billions of dollars on dry holes.

The young geologists possessed impeccable credentials. But their primary knowledge vastly exceeded their secondary knowledge—they overestimated what they knew. As they matured and saw which wells actually came in, geologists learned to "calibrate" their predictions. By the time geologists had many years of experience, Shell could trust that if they said wells on a dozen sites each had a 50 percent chance of producing oil, about half of the wells would produce. But "on the job training" was causing far more dry holes than accurately calibrated predictions would have yielded.

So Shell established a new kind of training program. Inexperienced geologists were shown summaries of past cases reflecting the myriad factors that might predict oil deposits. For each, they had to make a prediction with a numerical level of confidence. Then, they received feedback—they were told whether or not oil was actually found at each of the sites.

The training worked wonderfully. Now, when even junior Shell geologists predict there is a 50 percent chance of producing oil at a site, the company gets a hit about five times out of ten.

Training Weather Forecasters and Bankers

When people observe many similar past events, they often learn to estimate accurately the likelihood of future events of that kind.

Believe it or not, such regular feedback (plus training) has given weather forecasters an enviable record of reliability.

When weather forecasters make a forecast, within hours they see the actual outcome. Thus they learn to make good forecasts. If an experienced weather person says there is a 60 percent chance of rain, rain falls 59 percent of the time. When a 90 percent chance of showers is predicted, it rains 87 percent of the time. By requiring levels of confidence in weather prediction, supervisors in the National Weather Service have made the feedback more meaningful and created enormous benefits for themselves, for weather forecasters who wanted to improve, and for users of weather forecasts.

Studying the past can help bankers as well. George Moore, a former chairman of Citicorp, used to keep a list of his bank's ten biggest mistakes—loans that had cost it many millions. All credit officers were required to study the list. After he retired, however, the bank stopped using it. In his book, *The Banker's Life*, Moore commented: "If the officers responsible for Mexico and Brazil and Argentina . . . had known some of the past mistakes . . . they wouldn't have made the loans they made."

Most Professionals Can Learn to Calibrate Better

Professionals learn to calibrate when they receive timely, accurate feedback on their estimates. Groups that typically calibrate well include expert bridge players and experienced accountants.

In most professions, however, you won't get good feedback automatically. Thus you won't learn to calibrate unless you or your organization works to *create* feedback. We'll discuss systems for making sure you learn from the results of your decisions in Chapter 9. But you can obtain an important kind of feedback on your calibration in your own professional field without waiting for the outcome of the decisions you're making today.

Here's one useful approach: Create a test that asks for range

estimates and levels of confidence in response to questions relevant to your profession. Give this test to all professionals in your department and to all new employees.

A test for personnel managers might include such questions as:

1. Among professionals hired in the last five years, how many have failed to be promoted on schedule?
2. How many applicants for systems analyst positions did advertisements in major out-of-town newspapers generate?
3. Minorities and women comprise what percentage of our job applicants?

and so on.

Remember that *learning to recognize how much you know and don't know* is the key issue, not actually coming close to the right answers. If your personnel department recognizes that it doesn't know what share of professionals progress on schedule, it can do research to find out. But if personnel managers believe they know and they're wrong, they are likely to adopt incorrect policies without any research at all.

Simply give your best range estimates together with levels of confidence. Then secure the exact answers from appropriate sources. Does the frequency with which real answers fall within the ranges you provided correspond with the levels of confidence you gave? If the real answers fall within your ranges significantly *less* often than you anticipated, you fail the test.

Remaking a Loan Department

The results of these tests often astound people. (Your own score may be a shocker too.) But these tests can help transform decision-making in organizations.

One bank loan officer who heard us talk about overconfidence applied these principles. His bank's loan portfolio had been declining, but most of his colleagues were certain they knew their market. Their boss was sure his officers had no problems with overconfidence. His comment? "No one is more realistic than a banker."

To indulge his subordinate, however, the boss agreed to an overconfidence test. The boss and every single loan officer flunked. The test showed that every one of them misunderstood the lending limits and other crucial facts about their competitors.

The boss was shaken. He took immediate action. He required each loan officer to contribute information each week to what he labeled a "competitor alert file." He told each officer to check the file weekly until the department understood the competition more realistically.

Within a few weeks a memo inserted in the competitor alert file by a loan officer in a branch office showed that a major customer was considering a switch to a different bank. The other bank was not even one of the major area banks the boss had considered to be competitors. In fact, it was so small it couldn't legally meet the needs of the customer who was considering switching. But the smaller bank had formed a consortium with another small institution to seek the business legally.

Because the loan officer in charge of the account had been alerted, he managed to successfully head off the client's move. He saved about $160,000 in annual revenue. Good feedback on the accuracy of judgments can improve many business endeavors enormously. Yet most of us function in jobs with poor feedback on the quality of our own judgments and decisions.

Step Three: Ask Disconfirming Questions

The "confirmation bias" can lead intelligent people amazingly far astray from the truth. Once you've reoriented your estimating processes to provide accurate ranges and confidence levels rather than dangerous single-point predictions, you should next ensure that you regularly ask an appropriate number of **disconfirming questions.**

A penchant for disconfirming thinking can be a manager's best friend. At one small metalworking firm, for instance, an inventory audit had found a discrepancy between the actual amount of raw material at their plant and the amount that records showed should be there. Plant managers guessed that a thief had stolen the metal, since theft had been a problem in the recent past. The managers alerted their security force. They questioned employees. They even spent a weekend at the plant searching for clues and racking their brains for ideas about how the metal could have been stolen. The inventory system indicated that the amount of metal missing was the equivalent of a year's profits, and they were terrified of reporting it to their boss. But no thief could be found.

On Monday they dejectedly told the company's president of their problem. He asked only one question—but a disconfirming question: "How much raw material would have to have been stolen to account for the discrepancy?"

The question surprised the managers. They hadn't calculated the amount because it seemed irrelevant to the question of who the thief might be. After a few minutes of figuring, they replied, "A small mountain." The inventory system was reporting an impossible volume of metal missing.

The managers had worked so hard pursuing their theft theory that they had not asked basic questions that might test their premise: How much metal was "gone"? Would other explanations of the discrepancy make more sense than thievery? In a single stroke the president's question eliminated the possibility of theft: No one could have moved that many tons of metal off-site undetected. Within hours, they discovered that during a change in accountants, a procedure to update inventory records had been incorrectly applied, causing a huge accounting error.

We believe the president's ability to ask the right question was not just luck. He understood the art of disconfirming questions—questions seeking information that will show our initial judgment is wrong. Especially when numerous possibilities exist, it is a crucial skill to eliminate efficiently the many blind alleys one might get stuck in.

Most of us should work to develop that art. We need to judiciously seek a balance between evidence that will confirm and sharpen our theories, on the one hand, and looking for evidence that may disconfirm them on the other. In any intelligence-gathering process, we need to make a deliberate effort to ask disconfirming questions. Write down some potentially disconfirming questions to ask in each major information-gathering effort you undertake.

A few disconfirming questions might have changed the course of World War II. The entire German military machine failed to ask them after it discovered that Britain was learning the locations of its submarines in the early days of the war. The high command theorized (plausibly) that Britain obtained its intelligence from an extraordinary spy network. And a security campaign did uncover an alarming number of spies.

But the real reason for British success was that British logicians under Alan Turing had cracked Germany's codes. A simple disconfirming question ("Could they have broken our codes?") and an easily conducted test (a false report leaked to the British only in code, for example) might have shown the

Germans some really useful ways to interfere with British intelligence.

What Is Misleading You?

You should ask disconfirming questions not only about your ideas but also about your data sources. Often, estimates and forecasts are biased because a source of data is consistently providing wrong information.

For example, if you rely in hiring computer programmers on data provided by candidates themselves, you are likely to wind up with programmers who are better in self-promotion than in programming. The solution is to establish a policy of checking external data such as references and previous employers.

Often experiments designed to disconfirm your ideas are excellent investments. David Ogilvy, in his memoirs, says he succeeded in advertising because he was always ready to run a few ads he deemed to be losers. Invariably, some were big hits, leading him to revise his theories.

Step Four: Expose the Hidden Sources of Future Problems

Good managers try to "expect the unexpected." Of course the unexpected is difficult to anticipate—even though we know it is inevitable!

Thus, too many organizations fail to complete their projects within the time they originally forecast. Too many can't deliver products at the costs they originally hoped. And many get blindsided by nasty surprises and expensive errors.

Availability biases cause people to underestimate the difficulties they will face: Few of the unexpected events that *might* delay a project come easily to mind when people are planning it. And no matter how knowledgeable people may be, they retain the human tendency to underestimate the importance of information that is not readily available to them.

Successful professionals have discovered that a few standard techniques outlined below—the use of fault trees, scenarios, and prospective hindsight—can improve estimates dramatically. These techniques simply make people see the possible causes of problems more clearly.

Good planning techniques consume time in decision-making processes that are conducted under severe time pressure. Thus people are constantly tempted to omit them to "save time." But omitting them often causes unhappiness and losses. For example, in the mid-1960s Boeing engineers made projections of manufacturing costs for the 747 jumbo jet that were too low by nearly 50 percent. Thus the jet was priced far too low, and even after ten years Boeing still had not recovered its investment in this "winner."

In recent years companies such as Hewlett-Packard and Xerox have required efforts to anticipate problems in the early stages of their design processes, because they recognize that the alternative is frustration and high costs.

Fault Trees Reveal Obstacles

One of the simplest ways to make possible problems more available is to construct a fault tree, a diagram that details the many ways in which a complex system might fail. For example, if you wanted to examine the reasons why a restaurant might fail (so you could realistically assess and minimize the risk of investing in a new restaurant), you might construct a fault tree with the primary set of branches representing the major kinds of failures. Then you would list subgroups of causes within each major branch. In each subgroup, include an "all other" category. See the sample on page 109.

Fault trees help people see the possible causes of problems more clearly in their minds and therefore help them to consider realistically the likelihood of difficulties. To maximize a fault tree's usefulness, however, you must list problem areas on the tree as completely as possible. Otherwise, estimates will still suffer from an availability bias.

For example, one of us asked a group of eighty restaurant and hotel managers studying at the Cornell School of Hotel Administration to use a fault tree to estimate the likelihood of a number of causes of restaurant failures, including an "all other" category.

Within a main branch of the fault tree labeled "Decreasing Number of Customers," for instance, one group saw a list of six causes plus "all other." The listed causes included "incorrect pricing," "unclear image," "outdated restaurant concept," etc. The group was specifically warned to "remember to include the likelihood of any reason not specifically listed in the 'all other' category." But the managers listed an average of only 7 percent under "all other."

A second group saw a list of twelve causes under "Decreasing Number of Customers"—the original six plus six additional such as "poor quality of food," "poor quality of service," "changing competition," etc. The longer list of causes made the additional six causes much more available, and these managers assigned them a combined probability of 50 percent. The "all other" category decreased only to 4 percent. So the total probabilities assigned by the first group to its "all other" category should have been on the order of 54 percent. In other words, the first group vastly underestimated such causes of customer decline as poor food and poor service simply because they had not been listed on the original chart we had given them.

However, our research also showed that you can expect to reduce such availability biases dramatically by creating a fuller fault tree. We asked a third group of restaurant managers to estimate likelihoods on a fault tree that listed the original six causes and allowed space for the manager to list up to six more. Some managers who used this chart failed to list any additional causes. They estimated probabilities similar to those of the first, highly biased, group. But those managers who did write in six additional causes of their own choosing—which were often different from those we had listed for the second group—provided probabilities similar to those of the second group. We don't know for sure whether the probabilities produced by any of the managers were accurate. But we do know that thinking about a list of twelve causes would greatly benefit a manager trying realistically to avoid future problems.

Fault trees especially produce useful estimates if you break a large project into parts and try to estimate the time required for each component. You'll usually find that you can estimate some parts of the project accurately because you have done similar work before. You won't need a fault tree for these components. But after you have estimated the components based on your past experience, you can then create fault trees to gauge the problems you'll face in the parts of the project you know less well.

Using Scenarios

Your goal is not just to list systematically what could go wrong, but also to examine the *joint* effect of possible problems. This way you can understand the risks that lie ahead and possibly deal with likely problems early in the process when the costs of avoiding them are low.

Creating optimistic and pessimistic scenarios will help you explore possible interactions among causes and help you think about outcome patterns too complex for a fault tree to handle.

FAULT TREE FOR RESTAURANT

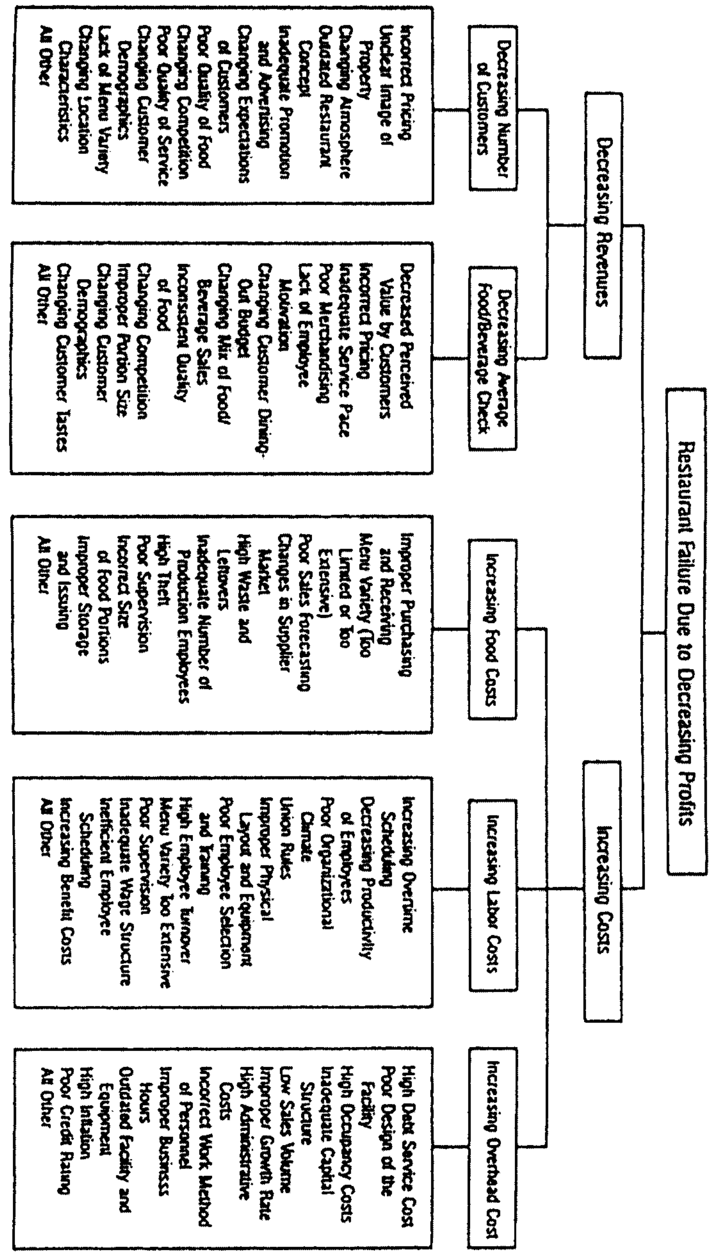

When you just *list* possible problems, you and your associates may tend to disbelieve them in isolation. It makes sense to discount the likelihood that a fire will destroy your factory. But when you try to draw even an optimistic scenario, you'll often recognize that all the developments you had wanted *could not* all happen. (For example, you may want to assume that a key person will be able to spend 90 percent of her time on the project. But you may also be assuming that your company will continue to receive enough new business in the period that no major budget cuts will be required, and receiving new business may imply that at least 40 percent of the key person's time will be spent meeting with new clients.)

Scenarios that include modest amounts of pessimism, moreover, may point out to you that certain combinations of reasons you listed in the fault tree could magnify each other rather than offset each other. For instance, an increased trade deficit may trigger an economic recession, which in turn increases unemployment and reduces domestic production.

Scenarios can help us understand how the various strands of a complex tapestry move when one thread is pulled. We've found that managers from almost any industry can create useful scenarios through the following process:

1. Define the major issues your project planning should consider. The issue definitions should mention each of the key variables involved in the project and the time frame over which each issue is important. (For example, one issue in an energy scenario might be, "Natural gas prices in the Eastern half of the United States over the next five years.")

2. List major players (including companies, governments, consumers, unions, and others). Think about each player's role, interests, and power position.

3. List potentially important economic, political, technological, and social trends that might affect the project.

4. Specify key uncertainties: Variables you cannot predict but that can impact your business positively or negatively (such as the level of the dollar).

5. Now construct two "forced" scenarios by putting all good outcomes in one and all bad outcomes in the other.

6. Assess the internal consistency and plausibility of these artificial scenarios. (For example, the "bad" scenario may include both the assumption that interest rates will be high and the assumption that the country will be mired in a deflationary recession. That's an unlikely combination.)

7. Rearrange the scenario elements to create at least two relatively consistent, plausible scenarios.

8. Evaluate the behavior of the major players in the scenarios created so far to determine whether the players' actions might change them. (For example, a competitor might undercut sales if you were able to increase prices as freely as your most optimistic scenario assumed.)

9. Distill all the information and insight you've gathered so far into a few (anywhere from two to four) distinctly different scenarios that cover a broad range and are each internally consistent.

In the notes to this chapter we have listed some sources helpful in scenario construction.

Prospective Hindsight

On the Monday after the football team has lost, everyone can see why the quarterback's pass to the tight end was a bad idea. Hundreds of "Monday-morning quarterbacks" declare that the quarterback should have anticipated the problem (since it's so easy for them to see it now). Whereas hindsight usually obstructs learning (as discussed in chapter 8), it can actually be turned to advantage in contemplating the future.

Experiments suggest that you can harness your ability to explain events so easily in hindsight to better anticipate *possibilities* that lie still ahead. Such prospective hindsight can be a valuable supplement to fault trees and scenarios.

Simply pretend that whatever you are evaluating has already occurred. For example, assume your project turned out to be a fiasco. Now try to list reasons *why* it occurred. You'll probably see *more reasons* for an event when you pretend it has already happened than when you simply ask yourself why it *might* occur. See the box on page 113 for an experiment you can try for yourself.

To test this technique, we used it with several groups of managers and MBA students. We gave each a brief description of a new employee including what kind of job, company, industry and personal ambitions. Using such vignettes, we asked half to generate plausible reasons why the new employee *might quit* six months from now. They generated an average of 3.5 reasons per person. The other half were told that the new employee *had quit* and were asked to generate plausible reasons why this might have occurred. The hindsight group generated 25 percent more reasons on average (a mean of 4.4 reasons). Moreover, the reasons were different in tone: more specific and linked to the episode presented.

We believe that effects due to prospective hindsight will even be stronger with real decisions involving large stakes. So if you doubt whether your people have sufficient insight into the myriad of causes that could produce success or failure for your project, ask them to engage in some "mental time travel."

PROSPECTIVE HINDSIGHT AS AN AID TO IMAGINATION

Test the effect of prospective hindsight on yourself and others. Create two groups. Give Question A to one group and Question B to the other.

A. How likely is it that a woman will be elected president of the U.S. in November 1996? Think about all the reasons why this might happen. For specificity, provide a numerical probability.

B. Imagine it is November 1996 and a woman has been elected president of the U.S. Think about all the reasons why this might have happened. Then provide a numerical probability of this actually occurring in 1996.

Most people find that Version B, in which the respondent pretends to be using hindsight, generates a greater number of paths to the event and the higher numerical probability.

Step Five: Limit Yourself to the Information You Can Handle

Lastly, ask whether you may be collecting too much information. If you are fully *rational*, more correct information can only help. However, for *normal* people, too much information can actually hurt, as the following study showed.

Paul Slovic and associates (then at the Oregon Research Institute) tested the predictive abilities of eight professional horse-race handicappers, racetrack employees who set the preliminary betting odds for each race. These handicappers must judge the likelihood that each horse will win the race. To do this, they use "past performance charts," which give nearly a hundred pieces of information on each horse and its history.

In the study, eight handicappers made predictions for numerous races. First, they were allowed to use only five pieces of information per horse from the charts—any five they wanted (often different pieces of information for different handicappers). Then they made the same predictions using ten pieces of information per horse. And then again with their most preferred twenty pieces, and one last time using forty pieces of information for each horse in the race.

Did the horse-race handicappers become more accurate in picking winners as they were given more information? They themselves felt that they would be. For each race, they were asked not only to pick the winner but also to estimate the chance that they were right. Their confidence increased as the information available rose.

But in reality, there was *no* significant improvement in their accuracy. Accuracy failed to rise, and their secondary knowledge—their realism about their predictive abilities—actually declined (see widening gap in Fig. 1 below).

We believe the additional information simply *overloaded* the handicapper, causing inconsistency and error. Other researchers have found similar results in other disciplines—for instance, with clinical psychologists who were asked to predict the behavior of patients with brain damage.

More information helps only to the extent that you can use it intelligently. Many professionals ask for too much in making estimates and decisions (to cover themselves or postpone having to make a decision). Numerous studies suggest that people have difficulty keeping more than seven or so "chunks" of information in mind at once. Vast amounts of data may only confuse matters, as the next chapter will demonstrate.

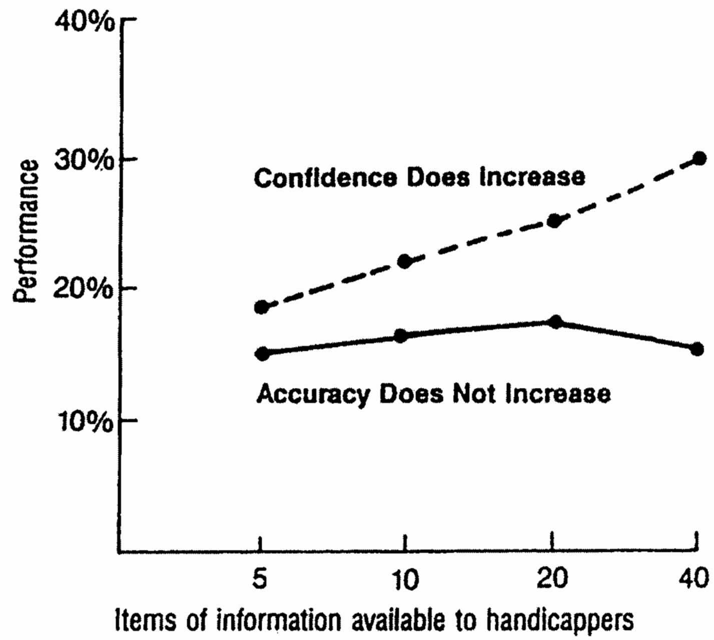

Figure 1
PREDICTIONS OF
HORSE-RACE HANDICAPPERS

Conclusions: How to Gather Intelligence Well

In this chapter, we've shown that a number of intelligence-gathering tactics can improve your success beyond what you can achieve simply by being aware of the dangers of overconfidence, the bias toward confirming evidence, the availability bias, and the anchoring trap.

A good decision-maker should:

—insist that people state estimates in terms of *ranges* and *levels of confidence*, so their meaning can be clear;

—create *feedback, training, and accountability* for people who make these estimates, so they can improve their skills;

—ask questions that might *disconfirm* initial presumptions (in case they are wrong);

—use techniques such as *fault trees, scenario building*, and *prospective hindsight* to seek reasons why projects may be more complex and difficult than you suspect at first;

—avoid collecting extensive amounts of information for estimates and decisions unless you have a way to synthesize it in a coherent fashion.

If you understand well the biases that all of us are prone to, and you take the above steps to counteract them, you'll create truly sound intelligence and lay a foundation for excellent decisions.