of these processes result in forgetting, which is a critical point of breakdown in human-system interaction. Techniques of system and task design and procedures to facilitate memory storage (training) were discussed.

In the next chapter, we discuss decision making, coupling the memory box in Figure 1.1 with the forward flow of information processing to include the selection of decision choices. Our treatment of decision making, however, depends on an understanding of memory and learning in three respects. First, many decisions place heavy loads on working memory. The costs imposed by these loads often lead to mental shortcuts, or **heuristics**, which produce systematic biases in decision performance. Second, other decisions are affected by long-term memory and experience. We decide upon an action because the circumstances correspond to a memory of a similar situation where we made the same decision, and that its outcome was successful. Finally, we will learn that the decision-making task has unique features, which cause learning and expertise in decision making to be somewhat different from that in other skills.

## Key Terms



# 8 DECISION MAKING

## 1. INTRODUCTION

Lauren had loved mountain climbing since she was a young girl, and in her twenties was now an accomplished climber. She decided to organize her own mini-expedition to climb a remote peak in the Northern Himalayas. To finance the expedition she took out a considerable loan on credit, and then turned to the choice of what mountain to tackle. There were so many options, varying in remoteness, altitude, challenge, uniqueness, beauty, possible weather, and information available. And then once the peak, Mt. Heuristic-Ri was chosen, the choice of team member was equally hard: how many, and whom? Friends she could trust or excellence of climbing reputation? And of her friends, good humor or strength, or organizational skills?

After a long trek in, they arrived at the foot of the mountain, but now were confronted by additional decisions: what route to take? What and how much equipment—was a higher camp necessary, or would they go for the summit in one long 24-hour shot—and what was the weather forecast? Unfortunately, it was rainy and cloudy for three days as they waited at base camp, until at last the weather begin to clear.

On the night prior to departure, the forecast, while iffy, indicated better weather on the way, so they decided to proceed with a 1:00 AM departure the next morning. Proceeding upward, the dawn was murky with clouds remaining over much of the sky, however a band of clear sky in west gave them hope and they continued onward. But the band never widened.

Leading the climb high on the mountain, Lauren was confronted with another choice above her: to veer to the left up a steep ridge of hard but solid rock, or to continue up an easier snow slope, burdened with new snow from the past several days of bad weather. The team was tired, and the snow looked good while the rock looked steep. Recalling her recent fall on a rock climb in Wyoming, Lauren chose the snow route. That choice almost proved disastrous; as the last climber neared the top of the slope, the large slab of snow below him started to cascade off. Fast action by the belayer above caught the climber before he was dragged down.

Topping the slope, they stopped to gaze at the sky, and noticed that the blue patch they had counted on was not opening, and indeed the ominous clouds behind them had grown. The summit was just about half mile beyond along the ridge, and Lauren huddled the team, saying "we are almost there. It might be risky to continue, but the summit is not far, and we have put so much into this expedition that we can't afford certain failure by turning back." Her teammate promptly rephrased the option: "If we turn back now, we'll surely get back safely, but if we continue there is still only a possibility that the weather will hold for us to make the summit." The team discussed the options briefly, and decided on descent as the safer course of action. While descending safely, Lauren remained somewhat dissatisfied. The weather had not turned worse, and she could only say "if only . . ."

Many serious accidents in which human error has been involved can be attributed to faulty operator decision making: The decision to launch the Challenger Space shuttle, which later exploded because cold temperature at launch time destroyed the seals is one example; another is the 1987 decision of personnel on board the USS Vincennes to fire on an unknown aircraft, which turned out to be a civilian Iranian transport rather than a hostile fighter (U.S. Navy, 1988). However a contrasting tragic decision was made by those on board the USS Stark cruising in the Mediterranean a year before, **not** to fire on an approaching target which turned out to be hostile and launched a missile which cost several lives on board the Stark. Of course these and other decisions gain notoriety because they generated unfortunate or tragic outcomes.

In the same manner we can recall better our own personal decisions that went awry: the class we chose to take that we failed; the poor investment we made, or Lauren's decision to take the snow slope. However in terms of frequency, our lives are far more dominated by the less salient decisions that went right. In this chapter we consider the processes that underlie decisions of both kinds, and the characteristics of the information and choice that can either improve the likely outcome, or make the decision more difficult and the choice more likely to produce an unwanted result.

## 2. CLASSES AND FEATURES OF DM

From an information processing perspective, decisions typically represent a many-to-few mapping of information to responses. That is, a lot of information is typically perceived and evaluated in order to produce a single choice. The following are some key features:

**Uncertainty**. An important feature of any decision is the degree of **uncertainty** of the consequences. Such uncertainty is generally a result of the probabilistic nature of the world in which we live, in which a given choice may lead to one sort of outcome if certain characteristics of the world are in effect or will come to pass, and a different outcome otherwise. If some of the possible but uncertain outcomes are unpleasant or costly ones, we usually consider the uncertainty of the decision as involving **risk**. The decision to purchase one of two possible vehicles is generally low risk if one has done advanced research on product quality, since the probable outcomes of one purchase or the other are known. But the decision to proceed with a flight in uncertain weather may have a high amount of risk, since it is difficult to predict in advance what impact the weather will have on the safety of the flight.

**Time**. Time plays at least two important roles in influencing the decision process. First, we may contrast "one shot" decisions like the choice of a purchase, with evolving decisions like those involved in treating an uncertain disease, in which test is followed by medication which may be followed by further tests and further treatment. Secondly, **time pressure** has a critical influence on the nature of the decision process (Svenson & Maule, 1993).

**Familiarity and Expertise.** Decision making changes with experience in several ways (Lipshitz & Cohen, 2005; Montgomery Lipshitz & Brenner, 2005; Weiss and Shanteau, 2003). As we discuss later, experts can often look at a decision problem and intuitively, nearly instantly pick the correct choice, whereas novices may ponder the problem for some time, and perhaps make a poor choice. This distinction parallels (although is not identical to) a dichotomy that research has distinguished between holistic and analytical decision types (Hammond et al., 1987), or between decision systems 1 (more holistic) and 2 (more analytic) (Evans, 2007; Kahneman & Klein, 2009; Kahneman, 2003; Sloman, 2002). Indeed these two systems appear to rely on different brain structures (Leher, 2010). In short, system 1 operates relatively automatically and effortlessly, reflecting "skilled expertise," and hence obviously develops fluency as the decision maker gains familiarity with a domain. System 2 is much more analytical and deliberative, generally relying heavily on working memory capacity in its deliberations. The two systems often work interactively, in that system 2 may monitor and cross check the quick intuitive decision made by system 1. As we will discuss below, the two are also somewhat associated with two different schools of decision analysis and research, **naturalistic decision making** (Zsambok & Klein, 1997, high skill and expertise: system 1) and the **heuristics/biases** approach to decision making (system 2; Kahneman & Klein, 2009).



**Classes of decision-making research.** Certain of the features of decision making described above have played a prominent role in distinguishing three important classes of decision-making research. The study of **rational** or **normative decision making** (e.g., Edwards, 1987) has focused its efforts on how people *should* make decisions according to some optimal framework; for example, one that will maximize the expected gain or minimize the expected loss. Efforts here are often focused on the departures of human decision making from these optimal prescriptions. We considered a simple example of this in the context of setting the "optimal beta" for signal detection decisions in Chapter 2 and will discuss it in more detail in Section 6 below.

The **cognitive** or **information processing** approach to decision making focuses more directly on the sorts of biases and processes that reflect limitations in human attention, working memory, or strategy choice, as well as focuses on common decision routines—known as **heuristics**—that work well most of the time, but occasionally lead to undesirable outcomes (Kahneman, Slovic, & Tversky, 1982; Herbert, 2010; Hogarth, 1987; Gilovich, Griffin, & Kahneman, 2002; Kahneman & Klein, 2009). Less emphasis here is placed on departures from optimal choice per se, and more on understanding the **causes** of such biases in terms of the structure and limits of the human as an information processing system. Finally, the naturalistic decision making approach (Kahneman & Klein, 2009; Mosier and Fischer, 2010; Zsambok & Klein, 1997, see Section 8) places its greatest emphasis on how people (usually experts) make decisions in naturalistic environments (i.e., outside of the laboratory), where they possess expertise in the domain and where the decisions have many of the aspects of complexity (evolving time, time pressure, multiple cues) that may be absent in laboratory studies of decision making (Mosier & Fischer, 2010).

## 3. AN INFORMATION PROCESSING MODEL OF DECISION MAKING

Figure 8.1 presents a model of the information processing components that are involved in decision making, elaborating the information processing presented in Chapter 1 while deemphasizing some components (e.g., sensory processing, response execution).

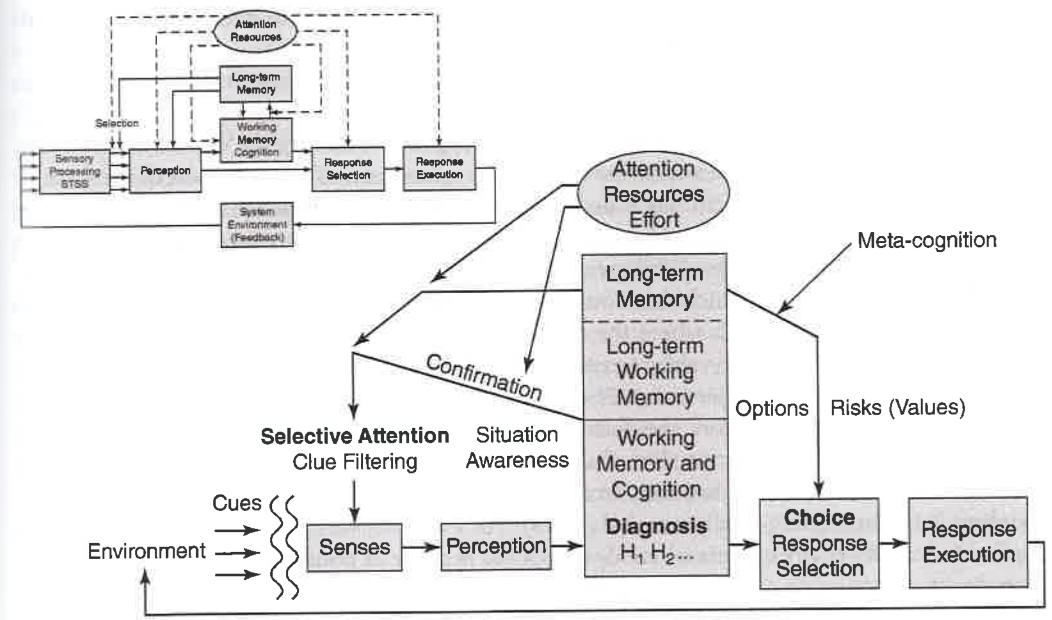


**FIGURE 8.1** An information Processing model of decision making. The general information processing model is shown in the upper left.

Beginning at the left, the decision maker seeks **cues** or information from the environment. However we note that in decision making (unlike much of pattern recognition), these cues are often processed through the "fuzzy haze" of uncertainty, and hence, may be ambiguous or interpreted incorrectly. In our opening story, Lauren was required to process the fuzzy uncertain weather forecast in making her decision to proceed. **Selective attention** of course plays a critical role in decision making, in choosing which cues to process (of higher perceived value) and which to filter out. Such selection is based on past experiences (long-term memory) and requires effort or attentional resources.

The cues that are then selected and perceived now form the basis of an understanding, awareness, or assessment of "the situation" confronting the decision maker (see Chapter 7), a process that is sometimes labeled **diagnosis** (Rasmussen & Rouse, 1981). Here the decision maker entertains hypotheses about what might be the current and future state of the world, upon which a decision should be based. For example, the physician must diagnose a disease before deciding upon a treatment, or the student may wish to assess an instructor's quality prior to choosing to enroll in a course. This diagnosis or assessment is based upon information provided from two sources, the external cues filtered by selective attention (bottom up processing) and long-term memory. The latter can offer the decision maker both various possible hypotheses of system state (e.g., the physician's knowledge of possible diseases and their associated symptoms or cures) and estimates of the likelihood or expectancy that each state might be true (top down processing). What makes decision making distinct from many other aspects of information processing is that diagnosis or situation assessment is often incorrect, because of the uncertain nature of the cues, their ambiguous mapping to possible states, or because of vulnerabilities in the cognitive processing of the decision maker related to selective attention (Chapter 3) and working memory (Chapter 7).

Many decisions are **iterative** in the sense that initial hypotheses will trigger the search for further information to either confirm or refute them. Troubleshooting a system failure will often trigger repeated tests to confirm or refute possible hypotheses (Hunt & Rouse, 1981). This characteristic defines the important feedback loop to cue filtering, labeled "confirmation" in Figure 8.1. The entire process of cue seeking and situation assessment has been labeled the "front end" of the decision process (Mosier & Fischer, 2010).

Following from the front end stages of cue seeking and situation assessment (or diagnosis), the third principle stage in decision making is the **choice** of an action, described as the "back end" of decision making (Mosier & Fischer, 2010). From long term memory the decision maker can generate a set of possible courses of action or decision options; but if the diagnosis of the state of the world is uncertain (as it is in much decision making), then the possible consequence of the different choices define their risks. Consideration of risk requires the explicit or implicit estimation of two quantities: (1) the *probability* or likelihood that different outcomes will come to pass and (2) *values*, the extent to which those outcomes are "good" or "bad." This is directly analogous to the discussion in Chapter 2, where the decisions made in signal detection theory depended upon both probability and the values (costs and benefits) imposed on different outcomes (hits, false alarms, misses, correct rejections). Thus the physician will probably consider the values and costs of various outcomes before she decides which treatment (do nothing, drugs, surgery) to recommend for a patient's abnormality of uncertain identity.

The overall distinction between front-end and back-end processes is critical to understanding decision failures (Hoffman et al., 1998). For example, very different solutions may be applied to remedy environments where decisions fail because of poor information and situation assessment, compared to those when failures result from inappropriate (e.g., too risky) choices in the face of a well-diagnosed situation (Wiegmann Goh & O'Hare, 2002).



Three additional components characterize our model. First, many of the components of decision making demand **effort or resources (see Chapter 10)**. Sometimes people choose decision strategies that impose reduced effort demands as they conserve this effort, such as choosing a diagnostic strategy that does not require them to hold many alternatives in working memory. Indeed such effort-conserving choices form a basis of many of the heuristics that we will discuss below.

Second, the figure depicts the role of **meta-cognition** (Reder, 1996). This process, discussed further in Section 7—awareness and knowledge of one's own knowledge, effort, and thought processes—is one that is closely linked with situation assessment (in this case, the "situation" involves the evolving decision process) and turns out to have an important influence on the overall quality of decision making: is one aware of the limitations in one's own decision process? Does the decision maker know that he does not possess all of the information necessary to make a good decision and hence seeks more?

Finally, the major feedback loop as shown at the bottom of Figure 8.1 critically illustrates the iterative nature of decision making. First, feedback of decision outcomes is sometimes used to assist in refining a diagnosis as we described above in troubleshooting. Second, meta-cognitive evaluation may trigger the search for more information. Third, feedback **may** be employed in a learning sense, to improve the quality of overall decisions (i.e., learning from one's mistakes); this feedback (although often delayed) may eventually be processed in long-term memory in order for the decision maker to revise his internal rules of decision making or the estimates of risks (see Section 8). That is, to learn decision-making skills.

## 4. WHAT IS "GOOD" DECISION MAKING?

The previous section has emphasized the several information processing components involved in decision making, such as cue perception, selective attention, and working memory. In previous chapters we have discussed many of these components in detail and have outlined some of the limitations (as well as the strengths) of all of them, such as the limited capacity of working memory. Hence, it is not surprising that the decision process may often fall short of "perfect" or "optimal" performance. Mistakes are made. Yet at the same time, the concept of what really is "good" decision making has proven to be illusive (Kahneman & Klein, 2009; Lipshitz, 1997; Shanteau, 1992), in contrast to other aspects of human performance, where speed and accuracy have a clearly defined status of quality. In fact, at least three different characterizations have been offered of "good" decision making, not all of which are in perfect agreement with each other.

First, early decision research of the normative school offered the **expected value** of a decision as the "gold standard." That is, the decision that would produce the maximum value if repeated numerous times (Edwards, 1987; see Section 6.1). However, defining expected value depends upon assigning universally agreed upon values to the various possible outcomes of a choice; values are often personal, making this a difficult undertaking. Even if values could be agreed upon, the choice that might be optimal if the decision is repeated time and again with plenty of time for weighing all the cues will not necessarily be optimal for a single choice, particularly one made under time pressure with little time to fully diagnose the situation and consider all possible outcomes (Zsambok & Klein, 1997). Furthermore, for a single decision, the decision maker may be more concerned about, say, minimizing the maximum loss (worst case) rather than maximizing expected long-term gain which after all can only be realized following a long-term average of the outcome of several decisions.

Second, one may say that "good" decisions are those that produce "good" outcomes and bad decisions conversely produce bad outcomes, such as the decision to launch the Challenger space shuttle, to fire on the Iranian Airliner from the USS Vincennes, Lauren's decision to climb the snow slope that triggered the avalanche, or the decision of a jury to convict a suspect who subsequently is found innocent. Yet we also know that in a probabilistic uncertain world, where cues are uncertain, it may only be in the 20-20 vision of hindsight that the decision can be labeled "bad" (Woods et al., 1994). After all, considering the USS Vincennes case (a "bad" decision), the decision makers on board the ship must also have considered that the decision made a year earlier on board the USS Stark, *not* to fire upon an approaching contact, turned out also to be "bad," leading to the loss of life on the Stark. This tendency to label a decision as good or bad only after the outcome is known is sometimes called the **hindsight bias**.

A third approach to decision quality has been based upon the concept of expertise (Zsambok & Klein, 1997, Kahneman & Klein, 2009; Brehmer, 1981; Shanteau, 1992, see Chapter 7). Since experts in other fields (e.g., chess, physics) are known to produce "good" and sometimes exceptional performance, why not consider that expert decision makers do the same. The problem here is that several analyses of decision making have shown that experts in certain domains do not necessarily make better decisions than novices (Brehmer, 1981; Dawes, 1979; Garling, 2009; Kahneman & Klein, 2009; Shanteau, 1992; Taleb, 2007; Tetlock, 2005; Serfaty, MacMillan, et al., 1997; see Section 8), and several "bad" decisions, according to our second criterion, have indeed been made by highly trained experts.

We adopt the approach here that, to the extent that all three of the characteristics described above converge then it becomes increasingly easy to discriminate good from bad decision making. But when they do not, then such discrimination is often fruitless, and it is much more appropriate simply to look at the qualitative ways in which different environmental and informational characteristics influence the nature of the processing operations and outcomes of the decision process. This is the framework shown in Figure 8.1, within which we treat the material below, first considering how people accumulate and assess evidence bearing on a diagnosis (front end: Section 5), then how they use that assessment to choose an action (back end; Section 6), and then the explicit role of effort and meta-cognition (Section 7).

## 5. DIAGNOSIS AND SITUATION ASSESSMENT IN DECISION MAKING

Accurate situation assessment is necessary (although not sufficient) for good decision making. Pilots who are good decision makers (by the various criteria above) actually take **longer** in understanding a situation or decision problem, even as they select and execute the choice more rapidly (Orasanu & Fischer, 1997). As shown in Figure 8.1, we can however distinguish four different information-processing components, each with their limitations, that can influence the quality of assessment and diagnosis: the role of **perception** in estimating a cue, the role of **attention** in selecting and integrating the information provided by the cues, the role of **long-term memory** in providing background knowledge to establish possible hypotheses or beliefs, and finally the role of **working memory** as the "workbench" for updating and revising beliefs or hypotheses on the basis of newly arriving information.

### 5.1 Estimating Cues: Perception

On the whole, people are reasonably accurate in estimating the mean and variance of a set of observations (Sniezek, 1980; Wickens & Hollands, 2000). However, systematic biases have been observed in perceiving and estimating three other characteristics of the environment: proportions, projections, and randomness.



**5.1.1** PROPORTIONS With regard to proportions, when perceiving a set of dichotomous observations (e.g., faulty versus normal parts on an inspection line; see Chapter 2), people do a reasonably accurate job of estimating the proportion so long as proportion values fall within the midrange of the scale (e.g., between around .05 and .95); however, with more extreme proportions, their estimates often tend to be "conservative," biased away from the extremes of 0 and 1.0 (Varey, Mellers, & Birnbaum, 1990). Such biases may result from an inherent conservative tendency ("never say never"), or alternatively they may result from the greater **salience**, noticeability or impact of the single outlying observation (which is, by definition, the infrequent event) in the sea of more frequent events. For example, if I have seen 99 normal parts, then detecting the one abnormal part will make more of an impact on my consciousness than detecting a 100th normal one. Its greater impact could well lead me to overestimate its **relative** frequency in hindsight, even as the rarity of the abnormal part will make me less likely to **detect** it in the first place if its abnormality is not salient (see Chapter 3).

However a very important exception to this overestimation bias occurs when the estimate of the frequency of *very* rare events (e.g., causing a rear end collision by following too close) is based on personal experience, rather than description (Hertwig & Erev, 2009). Here the person's sample of events is insufficient so they never actually experience the event in question, and underestimation is observed; that is, they may act as if the event is impossible (for them) rather than just improbable. This finding has important implications for safety as we discuss in Section 6.

The tendency to overestimate the frequency of rare events from description (versus experience, as above) has important implications for choice behavior. For example, people appear to show little difference in behavior (e.g., purchasing lottery tickets) whether odds of an event (winning) is 1/1000, or 1/10,000, and thereby implicitly overestimating the probability of the latter (Slovic & Finucane et al., 2002). They consider both as equal evidence for the **possibility** of winning rather than as different evidence for the **probability** of winning. In Chapter 2, we saw how this tendency could affect the setting of the response criterion, as manifest in a "sluggish beta." Later in this chapter, we see how it effects risky decision making.

**5.1.2** PROJECTIONS With regard to projection, humans are not always effective in **extrapolating non-linear trends**. As shown in Figure 8.2, they often bias their estimates toward the more linear extrapolation of the tangent where the data end (Waganaar & Sagaria, 1975; Wickens, 1992). This parallels the challenges people have in predicting the dynamic behavior of systems to be tracked, as discussed in Chapter 5. Thus, for example, in predicting the future temperature of a process on the basis of historical trend data of the exponential growth, people would be likely to underestimate its future values. Like the bias in estimating proportions, this can be thought of as a "conservative" one, inferring that the quantity will be less extreme than the statistical data would suggest. However, such prediction is, by definition, an inference, and so the conservative bias in extrapolation can possibly be explained on the basis of a further inference based upon past experience. This is the inference that most exponentially increasing quantities do eventually encounter self-correcting mechanisms that slow the rate of growth. For example, exponential population increases will encounter natural (through disease) or artificial means (i.e., of birth control) to lower the rate of growth. Exponentially increasing temperatures will often trigger fire extinguishing efforts, or opening pressure relief valves that will reduce the rate of growth. So the long-term memory of experience will lead the decision maker—accurately—to **infer** that the rapidly growing quantity will eventually slow its rate of growth.

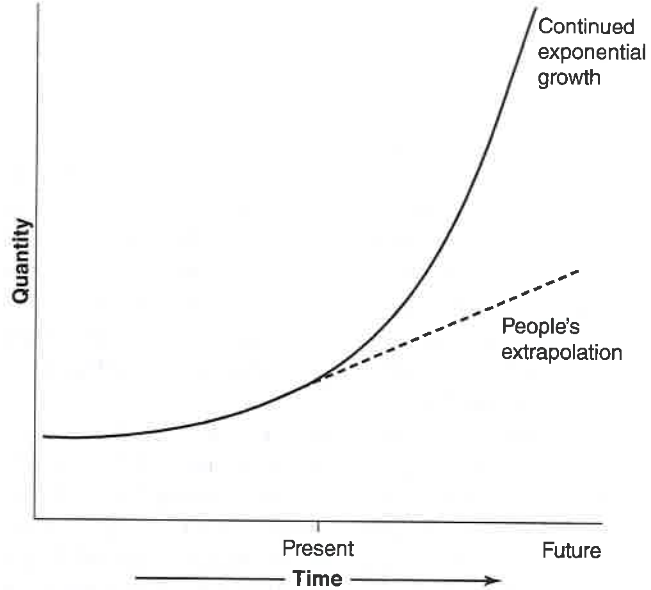


**FIGURE 8.2** Conservatism in extrapolation.

At the same time, other research indicates that people (e.g., stock analysts) may sometimes be overly risky or extreme in their projection of quantities that are *not* exponentially growing as above (De Bondt & Thaler, 2002), leading to an overreaction in their trading (e.g., choice) behavior. Indeed, they tend to be even more so when making longer-range forecasts, as if discounting the lower reliability of the greater look-ahead-time, a point to which we return below (De Bondt & Thaler, 2002, see also Chapter 5). Finally, as we discuss in Section 7.2, people are not always effective in planning for the future.

**5.1.3** RANDOMNESS People do not do a good job in perceiving (or understanding) randomness in the environment (Tversky & Kahneman, 1971). This is best illustrated by the **gambler's fallacy** in observing (or acting on) a series of dichotomous events, like coin tosses or wins and losses in a gamble. People tend to think that "random" implies a heavy bias toward alternation between the two outcomes. When generating a random series of say heads and tails, people will tend to avoid generating a sequence like HHH or TTTT, even as this sequence of three (or four) identical events is no less likely than any other sequence. In particular, when people witness a series of dichotomous events, the more consecutive observations of one event (e.g., losses) they see, the more they expect the next one to be the other event (a win). This is true despite the fact that in a random process, each event is independent of the prior one. The chances of a head following four heads is still 50 percent, not higher, as people's predictions would suggest.

This bias in the perception of random events is shown clearly in the "hot hand" effect in basketball (Gilvich, Vallone, & Tversky, 2002). Here, many players and coaches are convinced that after a few consecutive baskets (usually from outside) the player has a "hot hand" and should continue shooting (often at the expense of distributing the ball to teammates). Yet careful statistical analyses of such "streaks" indicate that they are no more likely than is the series of, say four "heads" in a coin toss. The next shot has a probability of success that is no greater than the player's



long-term shooting percentage. Indeed, if anything the streak could lead to the opponents' more aggressive defense against the hot hander, hence decreasing her likelihood of hitting the next shot.

Poor perception of randomness is also reflected in people's resistance to perceiving outliers in a distribution as legitimate components of the tails of an otherwise random distribution. They interpret them instead to be significant trends (Tversky & Kahnemann, 1981). People search for what they perceive to be systematic trends, and therefore they often see "patterns" in data which are, in fact, nothing more than random organization.

The previous discussion of biases in the perceptual estimation of quantities spawns one important design message. When possible, systems should display directly, the parameters estimated from separate observations (e.g., computer generated predictions), rather than requiring the human to estimate or infer those quantities. The format in which these parameters should be displayed (e.g., digital, graphical) was an issue discussed in earlier chapters of the book, and also has important implications for decision-making displays, as will be discussed toward the end of this chapter.

## 5.2 Evidence Accumulation. Selective Attention: Cue Seeking and Hypothesis Formation

As shown in Figure 8.3, we can represent the diagnostic stage of decision making as a process by which the decision maker receives a series of cues, symptoms, or sources of information as shown near the bottom, bearing on the true (or predicted) state of the world, and attends to

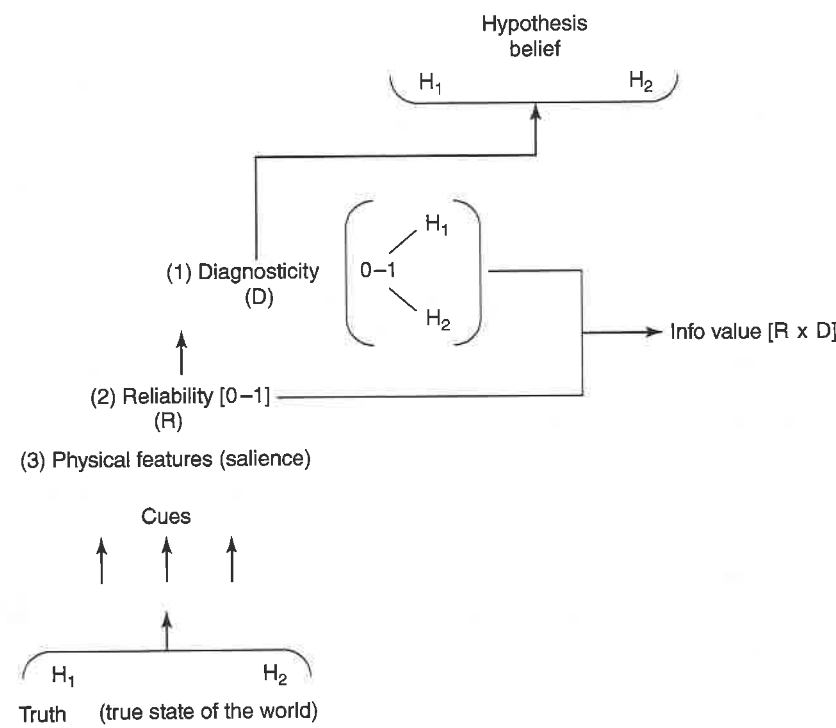


**FIGURE 8.3** Representation of the process of information integration (from bottom to top) to form a belief or diagnosis related to one of two hypotheses.

some or all of these with the goal of using those cues to influence the cognitive **belief** in one of several alternative hypotheses shown at the top right. In many instances, we can represent this as a "belief scale," between two alternative hypotheses, $H_1$ and $H_2$, as shown in the figure. Thus, we may think of the physician diagnosing a tumor as benign or malignant, the planner (for a flight, a hike, a picnic) predicting that the weather will be either clear or rainy, the investment broker predicting that the stock in a company will either climb or dive, or intelligence agents diagnosing the presence or absence of weapons of mass destruction in Iraq (Isakoff & Corn, 2006).

Each cue that potentially bears on the hypothesis can also be characterized by three important properties:

1. **Cue diagnosticity** formally refers to how much evidence a cue should offer regarding one or the other hypothesis. Thus, if one sees rain drops falling, this is a 100 percent diagnostic cue that it will be raining; on the other hand, a forecast of "a 50 percent chance of showers" is a totally undiagnostic cue for precipitation. Dark clouds on the horizon are relatively diagnostic (e.g., 75 percent), but not perfectly so. The diagnosticity of any cue can be expressed both in terms of its discriminating value (high or low) as well as its polarity (i.e., which hypothesis the cue favors).
2. **Cue reliability** or **credibility** refers to the likelihood that the physical cue can be believed. This feature is independent of diagnosticity. Thus an eyewitness to a crime may state categorically that "the suspect did it" (high diagnosticity); but if the witness is a notorious liar, his or her reliability is low. Collectively, both diagnosticity and reliability can be expressed on scales of 0 to 1.0, and then their **product** can reflect the **information value** of a cue. If the decision maker views a cue with an information value =1 ($d=1 \times r=1$), then that single cue is all that needs to be processed to make an error free diagnosis. However, most diagnostic problems have cues with information value less than 1.0, and hence can produce circumstances in which cues **conflict**. (Consider opposing witnesses for the defense and prosecution in a legal trial.)
3. The **physical features** of the cue which can make it conspicuous or **salient** have an important bearing on the selective attention and the subsequent processing that it receives.

How then should the multiple cues be integrated to form a belief that correlates with the true state of the world? Here we can consider four information-processing operations, three of them having parallels with our discussion of perception in earlier chapters. First, ***selective attention*** must be deployed to *process* the different cues, ideally giving different weight according to their information value. Second, the cue values—raw perceptual information—must be *integrated*, analogous to the *bottom up processing* of perceptual features in pattern recognition. Third, *expectancies* or prior beliefs may play a role in biasing one hypothesis or belief to be favored over the other, analogous to the way that expectancies stored in long term memory influence the *top down processing* in perceptual pattern recognition and signal detection (Chapters 6 and 2). Fourth, an operation that is not paralleled by those in perceptual pattern recognition, is the iterative **testing** and **retesting** of the initially formed belief, to attain the final belief which is the basis for choice.

Having established the role of reliability and diagnosticity in determining the information value of a cue, we are then in a position to establish the optimal degree of belief in one hypothesis or another on the basis of multiple cues.

The process of attending to and integrating multiple cues typically located at different places and/or delivered at different times along various sensory channels presents a major challenge to human selective attention and hence can be a source of four major vulnerabilities, as we discuss below.



**8.2.1 INFORMATION CUES ARE MISSING** A decision maker may not have all the information at hand to make an accurate diagnosis. An operator's judgment to turn on a faulty piece of equipment cannot be blamed if the operator was not informed by maintenance personnel of the equipment failure. At the same time however sometimes a decision maker may be blamed if a decision is made in the absence of critical information that she/he should know is essential. But thwarting this process is the fact that present cues can be **perceived**, while realizing the existence of absent cues depends upon **memory**, a process that we learned in the previous chapter is often prone to error. One quality of good decision makers is that they will often be aware (meta-cognition) of what they do not know (i.e., missing cues) and may proceed to seek these cues before making a firm diagnosis (Orasanu & Fischer, 1997). Thus, the effective planner of a mission will attempt to obtain, and rely on, only the most recent weather data, and if the available forecast is outdated may postpone a decision until a weather diagnosis can be made only on the most recent data.

**8.2.2 CUES ARE NUMEROUS: INFORMATION OVERLOAD** As we have noted, when the information value of any cue is known to be 1.0 (**both** reliability and diagnosticity = 1.0), then other information need not be sought. But this is rarely the case, and so effective diagnosis will rely upon integrating multiple cues. However, this can present a selective attention challenge, as we discussed in Chapter 3. The operators monitoring any nuclear plant in the face of a major failure may be confronted with literally hundreds of indicators, illuminated or flashing (Rubenstein & Mason, 1979). Which of these should then be attended first, as the operator then tries to form a diagnosis as to the nature of the fault.

When several different information sources are available, each with less-than-perfect information value, the likelihood of a correct diagnosis can increase as more cues are considered. In practice, however, as the number of sources grows beyond two, people generally do not use the greater information to make proportionately better, more accurate decisions (Allen, 1982; Dawes, 1979; Dawes & Corrigan, 1974; Lehrer, 2009; Malhotra, 1982; Schroeder & Benbassat, 1975). Oskamp (1965), for example, observed that when more information was provided to psychiatrists, their confidence in their clinical judgments increased but the accuracy of their judgments did not. Allen (1982) observed the same finding with weather forecasters. The limitations of human attention and working memory seem to be so great that an operator cannot easily integrate simultaneously the diagnostic impact of more than a few sources of information. In fact, Wright (1974) found that under time stress, decision-making performance deteriorated when more rather than less information was provided.

Despite these limitations, people have an unfortunate tendency to seek far more information than they can absorb adequately. The admiral or executive, for example, will demand "all the facts" (Samet, Weltman, & Davis, 1976). In the field of medical imaging, Jarvic et al. (2003) have noted that with the emergence of the MRI, surgeons begin to recommend a large number of unnecessary back surgeries, compared to the recommended rate when only lower quality X-rays were available. The extensively greater amount of information available in the MRI did **not** lead to improved diagnosis, and apparently degraded it (Lehrer, 2009).

To account for the finding that more information may not improve decision making, we must assume that the decision maker employs a selective filtering strategy to process informational cues. When few cues are initially presented, this filtering is unnecessary. When several sources are present, however, the filtering process is required, and it competes for the time (or other resources) available for the integration of information. Thus, more information leads to more time-consuming filtering at the expense of diagnostic quality.

**5.2.3** CUES ARE DIFFERENTIALLY SALIENT As we discussed with the SEEV model in Chapter 3, the **salience** of a cue, its attention-attracting properties or ease of processing, can influence the extent to which it will be attended and weighted in information integration (Payne, 1980). For example, loud sounds, bright lights, underlined or highlighted information, abrupt onsets of intensity or motion, and spatial positions in the front or top of a visual display are all examples of salient information cues and are likely to be given greater weight, particularly under time pressure (Wallsten & Barton, 1982). Negative, unpleasant information is found to be more salient (attention capturing) than positive, in influencing decisions (Yechiam, 2012).

These findings lead us to expect that in any diagnostic situation, the brightest flashing light or the meter that is largest, is located most centrally, or changes most rapidly will cause the operator to process its diagnostic information content over others: the **salience bias**. When integrating testimony from witnesses, it may be the loudest or most articulate voice that is attended to the best. It is important for a system designer to realize, therefore, that the goals of alerting (high salience) are not necessarily compatible with those of diagnosis in which salience should be directly related to the information value of the cue in making a diagnosis, not just in *detecting* a fault.

In contrast to salience, which may lead to "overprocessing," research also suggests that information that is difficult or effortful to interpret or integrate, because it requires arithmetic calculations or contains confusing language, will tend to be ignored, or at least underweighted (Bettman, Johnson, & Payne, 1990; Johnson, Payne, & Bettman, 1988). For example, Stone, Yates, and Parker (1997) found that presenting risk information in digital form led to less appropriate processing than presenting it in the analog form of stick figures, whose salient numerosity represented the magnitude of risk.

An extreme case of low salience relates to the **absence of a cue**. There are often circumstances in which a hypothesis can gain credibility on the basis of what is *not* seen as well as what is seen. For example, the computer or automotive troubleshooter may be able to eliminate one hypothesized cause of failure on the basis of a symptom that is NOT observed. Yet people are relatively poor in using the absence of cues to assist in diagnosis in fields such as medicine (Balla, 1980) or logical troubleshooting (Hunt & Rouse, 1981). It should be noted that the absence of a cue is not quite the same as the missing information described in 5.2.1 because there are circumstances in which the fact that something is NOT observed (absence of a cue) can provide a great deal of diagnostic information. It's just that people do not use that information very well.

The observation that cue salience influences the impact of cue processing is a part of the more general observation that the physical format or array of information relevant to a decision problem can influence the nature of the decision processes (Smith, Bennett & Stone, 2006), an issue we discuss in Section 7, and it also has relevance to the benefits of ecological interface displays with salient emergent features to the diagnosis of abnormal states in complex systems (Burns & Hajjckkk, 2008, see Chapter 4).

**5.2.4** PROCESSED CUES ARE NOT DIFFERENTIALLY WEIGHTED BY INFORMATION VALUE While people will tend to overprocess cues of greater salience, there is also good evidence that people tend to overprocess cues of lesser information value relative to those of greater value (e.g., Kohler, Brenner, & Griffin, 2002). That is, people do not effectively modulate the amount of weight given to a cue based upon its information value, whether the latter is influenced by diagnosticity *or* reliability. Instead, they tend to treat all cues *as if* they were more or less of equal value (Cavenaugh, Spooner, & Samet, 1973; Schum, 1975). This **as-if heuristic** thereby reduces the cognitive effort which would otherwise be required to consider differential weights when integrating information. It is a heuristic which, like others we discuss below, will not generally do damage to the



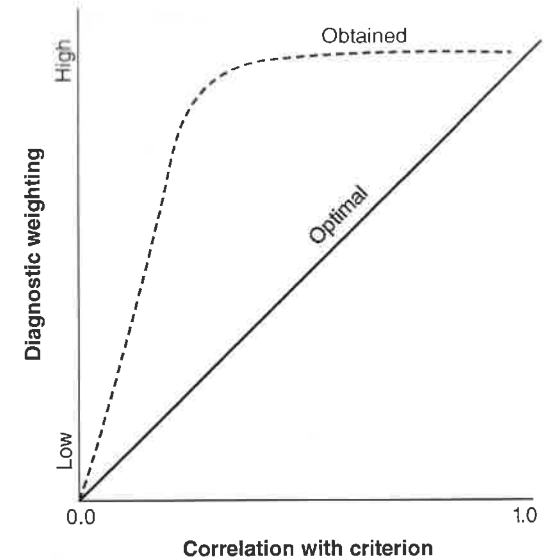


**FIGURE 8.4** Demonstration of the as-if heuristic. The function shows the relationship of the validity of cues to the optimal and obtained weighting of cues in prediction.

diagnosis (Dawes, 1979), but under certain circumstances, particularly when a low value cue happens to be quite salient, its use can invite a wrong diagnosis.

Kahneman and Tversky (1973) have demonstrated that even those well trained in statistical theory do not down-weight less reliable information sources when making "intuitive" predictions. In Figure 8.4, the optimal diagnostic weighting of a predictive variable is contrasted with the weights as inferred from subjects' predictive performance. Optimally, the information extracted, or how much weight is given to a cue, should vary as a linear function of the variable's correlation with the criterion. In fact, the weighting varies in more of an "all or none" fashion, as shown in the figure.

Numerous examples of the as-if heuristic can be identified, downweighting differences in information value. As one example, Griffin and Tversky (1992) found that evaluators, forming impressions of an applicant on the basis of letters of recommendations, tended to give more weight to the tone or enthusiasm of the letter (a **salient** feature) than to the credibility or reliability of the source (the letter writer). Koehler, Brenner, and Griffin (2002) found that when people make predictions, they generally neglect to consider differences in the quality of evidence, over-relying upon evidence when its quality is low, and under-relying when its quality is high. Rossi and Madden (1979) found that trained nurses were not influenced by the degree of diagnosticity of symptoms in their decision to call a physician. This decision was based only on the total number of symptoms observed.

A particularly dangerous situation occurs when less than perfectly informative information is passed from observer to observer. The lack of perfect reliability or diagnosticity may become lost as the information is transmitted, and what originated with uncertainty might end with certain conviction. There is some feeling, for example, that in the USS Vincennes incident in which the Iranian airliner was targeted, the uncertain status of the identity of the radar contact may have become lost as the fact of its presence was relayed up the chain of command (U.S. Navy, 1988).

Another potential cause of unreliable data whose limits are discounted in information integration occurs when the sample size of data used to draw an inference is small. A political poll based on 10 people is a far less reliable indicator of voter preferences than one based on 100. Yet these differences tend to be ignored by people when contrasting the evidence for a hypothesis provided by the two polls (Fischhoff & Bar-Hillel, 1984; Tversky & Kahneman, 1971, 1974). In the context of Figure 8.3, information regarding reliability can be said to be less **accessible** to cognition than the actual diagnostic content of that information, and hence is ignored (Kahneman, 2003).

The insensitivity to differences in predictive validity or cue reliability (e.g., optimal weighting) should make people ill suited for performing tasks in which diagnosis or prediction involves multiple cues of different information value. In fact, a large body of evidence (e.g., Dawes & Corrigan, 1974; Dawes, Faust, & Meehl, 1989; Kahneman & Tversky, 1973; Kleinmuntz, 1990; Meehl, 1954) does indeed suggest that humans, compared to machines, make relatively poor intuitive or clinical predictors. In these studies, subjects are given information about a number of attributes of a particular case. The attributes vary in their weights, and the subjects are asked to predict some criterion variable for the case at hand (e.g., the likelihood of success in a program or the diagnosis of a patient). Compared with even a crude statistical system that knows only the *polarity* of cue diagnositicity (e.g., higher test scores will predict higher criterion scores) and assumes equal weights for all variables, the human predicts relatively poorly. This observation has led Dawes, Faust, and Meehl (1989) to propose that the optimum role of the human in prediction should be to identify relevant predictor variables, determine how they should be measured and coded, and identify the polarity of their relationship to the criterion. At this point a computer-based statistical analysis should take over and be given the exclusive power to integrate information and derive the criterion value (Fischhoff, 2002).

Why do people demonstrate the as-if heuristic in prediction and diagnosis? The heuristic seems to be an example of cognitive simplification or effort conservation, in which the decision maker reduces the load imposed on working memory by treating all data sources as if they were of essentially equal reliability. Thus, a person avoids the differential weighting or mental multiplication across cue values that would be necessary to implement the most accurate diagnosis. When people are asked to estimate differences in reliability of a cue directly, they can clearly do so. However, when this estimate must be used as part of a larger mental aggregation using working memory, the values become distorted in this simplifying pattern.

## 5.3 Expectations in Diagnosis: The Role of Long-Term Memory

When cues are integrated, such integration is influenced in two important respects by long term memory (based on past experience), as related to cue correlation and to expectancy. Each generates its own unique heuristic.

**5.3.1 REPRESENTATIVENESS** The foundation of the **representativeness heuristic** (Kahneman & Frederick, 2002; Tversky & Kahneman, 1974) is that cues for a diagnostic state are often *correlated.* Thus, for example, bad weather is diagnosed by both clouds and low pressure. The flu is diagnosed by nausea, fever, and aches. The correlation between these cues or symptoms may be less than perfect. So there exists a difference between the ideal "prototype" (all cues present) and its actual expression in each real world "case." Some cues may be absent or weak, and possibly some extra cues may be present. When making a diagnosis, people tend to match the observed case pattern against one of a few possible patterns of symptoms (one for each diagnosis) learned from past experience and stored in long-term memory. If a match is made, that diagnosis is chosen. As we will see in Section 8, this is behavior typical of skilled decision making, or visual pattern recognition (Chapter 6).



There is nothing really wrong with following this heuristic *except* that people tend to use representativeness without adequately considering the **base rate**, probability, or likelihood that a given hypothesis or diagnosis might actually be observed (Koehler et al., 2002). For example, following the representativeness heuristic, a physician observing a patient who matches four out of five symptoms typical of disease X, and three out of five typical of disease Y will be likely to diagnose disease X as being most representative of the patient's symptoms, even if X occurs very rarely in the population, compared to disease Y.

In a manner similar to the failure to differentially weight cues, discussed above, Christenssen-Szalanski and Bushyhead (1981) have observed that physicians are insufficiently aware of disease **prevalence rates** (base rate) in making diagnostic decisions. Balla (1980, 1982) confirmed the limited use of prior probability information by both medical students and senior physicians in a series of elicited diagnoses of hypothetical patients. Furthermore, the sluggish beta adjustment in response to signal probability, described in Chapter 2, in which decision-making criteria are not adjusted sufficiently on the basis of signal frequency information, is another example of this failure to account for base-rate information. So too is the relative insensitivity to differences in proportion described in Section 5.1.1.

Representativeness may be thought to reflect another example of the distorting effects of salience or **accessibility** in decision making (Kahneman & Frederick, 2002; Kahneman, 2003). Symptoms are salient, accessible, and visible; probability is abstract and mental, and hence seems to be "discounted" when placed in competition with a pattern of perceivable symptoms. As Griffin and Tversky (1992) put it, "people pay more attention to the salient, representative strength of evidence (e.g., the difference between two means, or the warmth of description of an applicant in a letter) than they do to the reliability of evidence" (which is more abstract).

The prevalence of the representativeness heuristic does not mean that people ignore probability or base rates altogether in reaching diagnoses. It only means that physical similarity of expressed cues to a prototype hypothesis *dominates* probability consideration when the two are integrated to determine the most likely hypothesis, on the basis of both past experience and the physical evidence (Griffin & Tversky, 1992). If, on the other hand, the physical evidence is itself ambiguous (or missing), then people *will* use probability. They will be quite likely to diagnose the hypothesis which, *in their mind*, has the greatest probability of being true (Fischhoff & Bar-Hillel, 1984). However, this mental representation of probability may also be imperfect, as reflected in the second important heuristic in evidence consideration, the **availability heuristic**.

**5.3.2 THE AVAILABILITY HEURISTIC** Availability refers to "the ease with which instances or occurrences [of a hypothesis] can be brought to mind" (Tversky & Kahneman, 1974; Schwarz & Vaughn, 2002) and is closely related to the construct of accessibility discussed briefly above (Kahneman, 2003; Kahneman & Frederick, 2002). This heuristic can be employed as a convenient means of approximating prior probability, in that more frequently experienced events or conditions in the world generally **are** recalled more easily. Therefore, people typically entertain more available hypotheses.

Unfortunately, other factors strongly influence the availability of a hypothesis that may be quite unrelated to their absolute frequency or prior probability. As we noted in our discussion of long-term memory (Chapter 7), recency is one such factor. An operator trying to diagnose a malfunction may have encountered a possible cause recently, either in a true situation, in training, or in a description just studied in an operating manual. This recency factor makes the particular hypothesis or cause more available to memory retrieval, and thus it may be the first one to be considered. Lauren's recent fall on the rock in Wyoming led her to diagnose the rock route as more dangerous.

Availability also may be influenced by hypothesis *simplicity*. For example, a hypothesis that is easy to represent in memory (e.g., a single failure) will be entertained more easily than one that places greater demands on working memory (a compound double failure). Another factor influencing availability is the elaboration in memory of the past experience of the event. For example, in an experiment simulating the job of an emergency service dispatcher, Fontenelle (1983) found that those emergencies that were described in greater detail to the dispatcher were recalled as having occurred with greater frequency.

Availability and accessibility are closely related to the phenomenon of a **attribute substitution** (Kahneman, 2003) in which certain highly accessible mechanisms get substituted by the intuitive (type 1) decision system, for more effort-demanding mechanisms employed by the analytic (type 2) system when resources are scarce. Thus, highly accessible attributes like similarity, averages, and change are contrasted with (and often substitute for) more abstract, less accessible, but often more optimal attributes such as likelihood (influenced by probability) and absolute amount. As one simple example, when people make choices in a gamble, they are often heavily influenced by the probability of winning or losing between two options, rather than the expected value of the two options (an issue that will be discussed later in the chapter). Probability is bounded (by 0 and 1) and is easily accessible, comparable or discriminable between them (Slovic, Finucane, et al., 2002).

Interestingly, representativeness (the pattern of data) and availability (estimating frequency of hypothesis) are two commodities that are integrated together in the **Bayesian** approach to optimal decision making (Edwards, Lindman, & Savage, 1963). In this approach, prior probability is multiplied by the P(data pattern/hypothesis) to estimate the true probability of each hypothesis given the data. The interplay between availability and representativeness in human cognition approximates this process, as we saw too in signal detection judgments. In contrast, however, classical statistic fails to consider the prior probability (odds), focusing only on the "p value" or p(data/hypothesis). As we see by considering representativeness and availability, the human as an "intuitive statistician" considers both, but does so heuristically.

## 5.4 Belief Changes Over Time

As we have noted, many diagnoses are not the short, "one shot" pattern classifications, but rather take place over time as an initial tentative hypothesis may be formed, and more evidence is sought (or arrives) to confirm or refute it. Indeed most troubleshooting seems to work this way, in which various tests are performed, specifically designed to provide new cues or evidence in an effort to identify the "true" state. Jurors in a criminal trial also may form an initial hypothesis or degree of belief in the guilt or innocence of the suspect, but find these beliefs altered as further evidence is presented. Scientists form hypotheses and then design experiments and use subsequent data to either strengthen or weaken (usually the former; see 5.4.2) their belief in the hypothesis. In this process of refining beliefs over time, we can identify two important characteristics that can sometimes work against the most accurate estimate of the "truth": the anchoring heuristic and the confirmation bias. Later in the chapter we will also show how the *overconfidence bias* amplifies these two influences.

**5.4.1** ANCHORING HEURISTIC The **anchoring heuristic** (Einhorn & Hogarth, 1982; Chapman & Johnson, 2002; Joslyn et al., 2011; Kahneman & Tversky, 1973; Mosier, Sethi, et al., 2007) describes how, when cues bearing on a hypothesis, or information sources bearing on a belief arrive



over time, the initially chosen hypothesis tends to be favored, as if we have attached a "mental anchor" to that hypothesis and do not easily shift it away to the alternative. If evidence *a* favors hypothesis A and *b* favors B, then receiving the evidence in the order $a \rightarrow b$ should lead to a favoring of A, but receiving it in the order $b \rightarrow a$ will favor B. Such a tendency is consistent with the general observation that "first impressions are lasting."

One clear implication of the anchoring heuristic is that the strength of belief in one hypothesis over another will be different, and may even reverse depending on the *order* in which evidence is perceived (Adelman et al., 1996; Hogarth & Einhorn, 1992; Ricchiute, 1998). Allen (1982) has observed such reversals as weather forecasters study meteorological data on the probability of precipitation, and Einhorn and Hogarth (1982) have considered similar reversals as people hear evidence that is either supporting or damaging to a particular hypothesis about an event, such as jurors hearing different pieces of evidence for the guilt or innocence of a suspect (Ruva & McElvoy, 2008; Kahneman & Klein, 2009).

It should be noted that while anchoring represents a sort of **primacy** in memory (see Chapter 7), there is also sometimes a **recency** effect in cue integration, in that the most recently encountered of a set of cues may, temporarily, have a strong weighting on the diagnosis (Rieskamp, 2006). Thus the lawyer who "goes second" in presenting closing arguments to a jury may well leave the jury with a bias toward that side, in making their judgment of guilt or innocence (Davis, 1984).

Indeed, a careful review of studies and a program of experiments carried out by Hogarth and Einhorn (1992) revealed that a number of factors tend to moderate the extent to which primacy (anchoring) versus recency is observed when integrating information for a diagnosis. For example, primacy is dominant when information sources are fairly simple (e.g., a numerical cue rather than a page of an intelligence report), and the integration procedure is one that calls for a single judgment of belief after receiving all of the evidence, rather than a revision of belief after each piece of evidence. However, to the extent that the sources are more complex and hence often require an explicit updating after each source is considered, then recency tends to be more likely.

To add to the complexity of this analysis, a case can be made that in many dynamic circumstances recency *is* in fact more optimal (and anchoring less so) to the extent that the reliability of a given piece of sampled information declines over time. Thus in a sequence of patient health status reports, those encountered first, perhaps several hours old, should be somewhat discounted. Yet people do not do much of this age-related discounting (Wickens, Ketels, et al. 2010), still showing primacy and anchoring.

Whether primacy or recency is observed, in arguing for such innovations as integrated graphics displays for decision support (Bettman, Payne, & Staelin, 1986; Cook & Smallman, 2008; MacGregor & Slovic, 1986; see also Chapter 12) or simultaneous displays of unit/price information of a number of comparable products (Russo, 1977), researchers have made a convincing case that where possible, evidence that is available simultaneously should be presented simultaneously and not sequentially (Einhorn & Hogarth, 1981). A simultaneous format cannot guarantee that simultaneous processing will occur, which of course depends on the breadth of attention and the operator's own processing strategies. At least, however, it gives the operator the option of dealing with the information in parallel if attentional capabilities allow or of alternating between and revisiting different information sources, if they do not. In this manner, one information source is not given automatic primacy (or recency) over others.

**5.4.2** THE CONFIRMATION BIAS Evidence bearing on a hypothesis or belief may be either passively received or actively sought. The **confirmation bias** describes a tendency for people to seek information and cues that *confirm* the tentatively held hypothesis or belief, and not seek (or discount) those that support an opposite conclusion or belief. Ambiguous cues, that information which is totally undiagnostic within the framework presented in Section 5.1, will be interpreted in a manner that supports the favored belief (Cook & Smallman, 2008; Einhorn & Hogarth, 1978; Herbert, 2010; Hope, Memon, & George, 2004; Mynatt, Doherty, & Tweney, 1977; Nickerson, 1998; Schustack & Sternberg, 1981). This bias produces a sort of "cognitive tunnel vision" in which operators fail to encode or process information that is contradictory to or inconsistent with the initially formulated hypothesis, hence conferring even greater rigidity to the anchor.

The investigation into the USS Vincennes incident in the Persian Gulf revealed the confirmation bias at work. Operators of the radar system hypothesized early on that the approaching aircraft was hostile, and they did not interpret the contradictory (and as it turned out, correct) evidence offered by the radar system about the aircraft's neutral status (U.S. Navy, 1988). The analysis of the Three Mile Island incident also reveals a confirmation bias for the operators to confirm their belief in the erroneous hypothesis of a high-water level in the reactor (Rubenstein & Mason, 1979).

Arkes and Harkness (1980) demonstrated the selective biasing of *memory* induced by the confirmation bias. They presented subjects with several symptoms related to a particular clinical abnormality (experiment 1) or to the state of a hydraulic system (experiment 2). Arkes and Harkness found that if the subject held a hypothesis or made a positive diagnosis, the symptoms they had observed that were consistent with that diagnosis were readily remembered, whereas inconsistent symptoms were more easily forgotten. Furthermore, subjects erroneously reported seeing symptoms that they actually had not seen but that were consistent with the diagnosis. Similar observations of false memories for consistent cues were made in a study of aviation fault diagnosis by Mosier, Skitka, et al. (1998).

In a comprehensive review of the confirmation bias, Nickerson (1998) identified several possible reasons for this failure to seek disconfirmatory evidence:

1. People have less cognitive difficulty dealing with positive information than with negative information (Clark & Chase, 1972, see Chapter 6), and with the presence of information (a present cue that supports what you already believe) than the absence (the absence of a cue which, if present would support your belief), also reflecting cognitive effort. The process required to change hypotheses—abandon an old one and reformulate a new one—requires a higher degree of effort than does the repeated acquisition of information consistent with an old hypothesis (Einhorn & Hogarth, 1981). Given a certain "cost of thinking" (Shugan, 1980) and the tendency of operators, particularly when under stress, to avoid troubleshooting strategies that impose a heavy workload on limited cognitive resources (Rasmussen, 1981), operators tend to retain an old hypothesis rather than go to the trouble of formulating a new one, or even entertaining two hypotheses at one time, so long as accepting "the chosen one" is consistent with most of the evidence (e.g., close to the truth).
2. There is a motivational factor related to the desire to believe. The high value that people place on *consistency* of evidence leads them to see all (or most) evidence supporting one or the other belief, and that belief is usually the one initially formulated.
3. A second motivational factor results when people focus more on the consequences of the logical **choice of action** that would follow from the initially favored hypothesis, rather than the truth of that hypothesis itself (Bastardi, Uhlman, & Ross, 2011). As we will see below, choices are inherently value laden, given the likelihood of positive and negative outcomes that can flow from those choices in an uncertain world. Lauren was inclined to believe that the weather would clear because the consequences would be summiting and success of the expedition. Hence people may be inclined to stick with (and try to confirm) the belief supporting choices whose outcomes, if the belief is true, are less negative and more positive. As Nickerson says, "when using a truth seeking strategy [trying to disconfirm] would require taking a perceived risk, survival is likely to take precedence over truth finding." Often, finding one's beliefs to be wrong can be embarrassing.
4. In some instances it may be possible for operators to influence the outcome of actions taken on the basis of the diagnosis, which will increase their belief that the diagnosis was correct. This is the idea of the "self-fulfilling prophecy" (Einhorn & Hogarth, 1978). It might describe a teacher who, diagnosing a child as "gifted," will provide that child with sufficient extra opportunities and motivation so that high academic performance will be almost guaranteed. It might also describe the scientist who, believing a theory to be correct, will now design and carry out experiments that are most likely to produce confirming evidence.



The issue is how to force a diagnostician simultaneously to entertain alternative hypotheses and to seek disconfirming evidence or at least attend to it if it arrives—in short, to break through the cognitive tunnel. This represents a major challenge to the designer of systems in which troubleshooting will be required.

Finally, we note in the context of both the confirmation bias and anchoring, the insidious role of the **overconfidence bias** in amplifying the distorting influence of both. While this bias will be discussed in detail later in Section 7.2, for now we consider that to the extent that people are more confident than they have a right to be in their existing beliefs, then they will be even less likely to seek evidence that those beliefs may be wrong, creating a sort of vicious cycle or "perfect storm" of these biases. This scenario was played out in the conviction that Iraq possessed weapons of mass destruction, leading up to the Iraq war (Isakoff & Corn, 2006).

**5.4.3** DECISION FATIGUE A third influence on decision making over time is known as **decision fatigue** (Tierney, 2011). Repeated decisions can often lead to decreased effort invested in accuracy and analysis. This phenomenon was illustrated dramatically in an analysis of parole board decisions carried out by Danzigera, Levav, and Pesso (2011), who observed that the probability of granting parole declined from 75 percent early in the morning, down to approximately 25 percent later in the day. Stated simply, the effort or cognitive resources required to make careful decision analysis was depleted over time, such that the "effort-lite" default strategy of denying parole (essentially deciding not to decide) begins to dominate.

## 5.5 Implications of Biases and Heuristics in Diagnoses

The previous sections may have painted a fairly pessimistic picture of the accuracy of the human as a diagnostician, full of biases and heuristics that force beliefs away from "the truth." Although such departures are often observed and records are replete with examples of incorrect diagnoses (jury verdicts that have later been found incorrect; Three Mile Island, USS Vincennes, misdiagnosed diseases), several qualifications need to be applied to the view that humans are just "a bundle of biases." First, as we noted above, many of the heuristics are highly adaptive, for a decision maker who must work rapidly and cannot afford to invest a large amount of mental effort (and/or time) to consider all the symptoms and all possible hypotheses (Payne, Bettman, & Johnson, 1993). Indeed, heuristics are so often used by people precisely because *most* of the time they *do* provide a correct or at least satisfactory outcome (Gigerenzer et al, 2002; Gigorenza, 2002). If they were wrong more often than right, people would eventually abandon them (although see Section 8 below). Secondly, using the shortcuts offered by heuristics often is a necessity given the time constraints of a decision environment. For example, the fire captain must depend upon the speed of the representativeness heuristic in certain time-critical situations, when a delay in selecting an action can result in loss of life. And the confirmation bias can at times provide a very useful and adaptive way of gathering information (Klayman & Ha, 1987).

Finally, for all of the biases and heuristics described above, decision research has examined certain conditions under which they may be modulated or eliminated entirely. For example, overconfidence in forecasting appears to be eliminated from the forecasts offered by meteorologists (Murphy & Winkler, 1984; but not by experts in many other professions, Shanteau, 1992; see Section 8). Anchoring may be reduced or eliminated by properties of the cues (Hogarth & Einhorn, 1992). And there are great differences between circumstances and people in the amount of overconfidence in diagnostic estimates (Paese & Sniezek, 1991). What is most critical from the perspective of this book is that analysis of these sorts of biases can lead to suggested training, procedural, and design remediations which can lessen their degrading impact on diagnosis in the circumstances when those impacts may be severe, or safety compromising. We discuss these remediations in the final section of this chapter.

## 6. CHOICE OF ACTION

Up to this point our discussion of decision making has focused on a collection of processes involved with estimating the state of the world and diagnosing or making a situation assessment. These processes are necessary to sustain effective decision making, but are not sufficient. As represented in Figure 8.1, the output of decision making must also include a choice of some action. In this regard, the dichotomy of state assessment and action choice is analogous to that discussed in Chapter 2 (signal detection), between the evidence variable (representing the likelihood of a signal), and the response criterion (by which the evidence variable was transformed into a dichotomous choice). Lauren, our climbing leader, assessed the difficulty and safety of rock versus snow, and then chose the snow course of action.

One key feature of this choice, is not relevant for diagnosis but was clearly represented by signal detection theory is the **value** that the decision maker places on different possible outcomes. We consider below, how people "should" and how they do combine information on value and probability to make decisions, just as, in our discussion of signal detection theory, we considered how they combined information on values and probabilities in setting beta for the decision of whether a signal was present or not. We discuss first the nature of decisions that consider values only; then we consider the added complexity of combining probability with value when examining decision making *under uncertainty.*

### 6.1 Certain Choice

When choosing which product to buy, or Lauren's choice of teammates for the expedition, the choice can be often be conceptualized as in Figure 8.5, in which an array of possible **objects** (e.g., products) are compared, each with varying **attributes**. For example the set of personal computers to purchase may vary in their attributes of price, usability, maintainability, warrantee, and so forth. In making such a choice that will maximize the consumer's overall satisfaction, the decision maker should carry out the following steps:

1. Rank order the **importance** of each attribute (highest number, greatest importance). In Figure 8.5, the left attribute (price) is least important (1), the next attribute across (warrantee) is number 4 and so forth.



| | Attributes: Price | Warrantee | • • • • • • • |
|---|---|---|---|
| Importance: | 1 | 4 | |
| Objects A | 2 | 3 | |
| B | 3 | 1 | |
| • • • | | | |

Object A: $2 \times 1 + 3 \times 4 = 14$
Object B: $3 \times 1 + 1 \times 4 = 7$

**FIGURE 8.5** Choice under certainty. The calculations at the bottom are based on a choice between only two objects, although the extended rows and columns suggest that the procedure could generalize to many more objects and attributes.

2. Assess the **value** of each object on each attribute (highest number, greatest value). For example, the highest number would be for the least expensive product, the best warrantee, etc.
3. For each object assess the sum of the products of (value x importance), as is shown in the bottom of the figure.
4. Chose to purchase the object with the highest sum of products. As the calculation shows, in the example of Figure 8.5, this turns out to be object A.

This decision process is known as a **compensatory** one, in that a product which may be low on the most important attribute (an expensive computer, when cost is most important), can still be chosen if this deficiency is *compensated* for by high values on many other attributes of lesser importance. For example the most expensive computer may have far and away the best user interface, the most reliable maintenance record, and the best warrantee, allowing these strengths to compensate for the weakness in price.

While people may, in the long run, best satisfy their own expressed values by following the prescriptions of the compensatory method, many choices in everyday life are made with much less systematic analysis, following heuristics or other shortcuts (Leher, 2010). For example the rule of **satisficing** (Simon, 1955) is one in which the decision maker does not go through the mental work to chose the best option, but rather one that is "good enough" (Lehto, 1997), and this is often the strategy employed in real-world naturalistic decision making, when there is time pressure (Klein, 1989, 1997; Mosier & Fischer, 2010).

A more systematic heuristic that people sometimes employ when the number of attributes and objects is quite large, is known as **elimination-by-aspects** (EBA; Tversky, 1972). Here, for example, the most important attribute is first chosen, then any product that does not lie within the top few along this attribute (aspect) is eliminated from consideration, and then the remaining products are evaluated by comparing more of the aspects of the remaining few objects. As a heuristic, this technique will easily reduce the cognitive effort of needing to compare all attributes across all objects. And it will usually prove satisfactory, only failing to pick a satisfactory choice if an object that is low on the most important attribute (and hence eliminated) happens to be near the top on all others. Understandably, the EBA heuristic is one that begins to dominate over time, as people suffer the effort depletion of decision fatigue (Tierney, 2011).

## 6.2 Choice Under Uncertainty: The Expected Value Model

Unlike those choices discussed in the previous section in which the consequences of the choice were relatively well known, many decisions are made in the face of uncertainty regarding their future consequences. Such uncertainty may result because we do not know the current state of the world; for example a physician may choose a particular treatment, but be uncertain about the diagnosis. Lauren was uncertain of the avalanche conditions of the snow route. Others may result because the future cannot be foretold with certainty. Stock brokers are certainly vulnerable to accurately predicting the future market forces, prior to making investment decisions (De Bondt & Thaler, 2002; Kahneman & Klein, 2009; Taleb, 2007).

Indeed we can often represent decision making under uncertainty as shown in Figure 8.6, by providing the possible states of the world (A, B, C, . . .) across the top of a matrix, each associated with their estimated probability or likelihood, and the possible decision options (1, 2, . . .) down the rows. The representation in Figure 8.6 echoes three other analyses considered earlier. First, the estimated probabilities of states of the world, can be thought of as being "passed on" from the degree of belief in one of two or more hypotheses, as represented in Figure 8.3 and now shown at the top of Figure 8.6. Second, the matrix shares an analogous form with the *certain choice* matrix shown in Figure 8.5, and indeed the computations for the optimal choice are similar between the two matrices. Third, the matrix is in fact a direct analog to the signal detection theory decision matrix discussed in Chapter 2, with its two states of the world and two choices. However, in the context of the present chapter there may be more than two states of the world and more than two decision options.

As you will recall, a key aspect of the discussion of signal detection theory was the setting of the optimal beta, in a formula that was determined by the probability of the two states of the world, and by the *outcome costs and values* of the different states of the world that would be forecast from the four joint events. In Figure 8.6, these costs and values are represented by a value (V) (which can be either positive or negative) of the outcome associated with the consequence of each decision option made in each state of the world. One might consider for example the costs

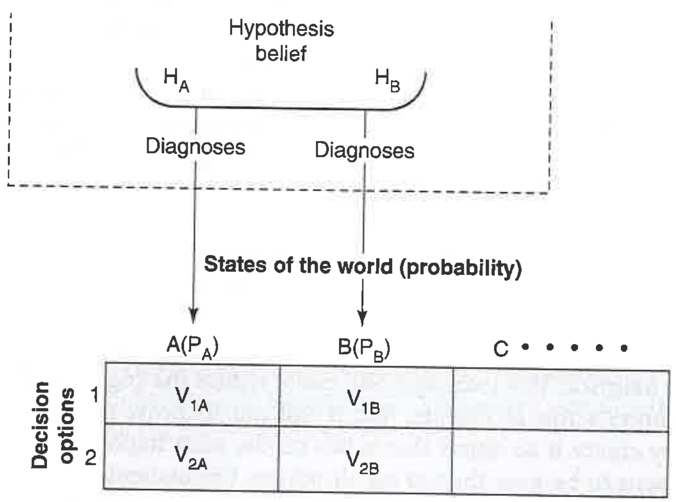


**FIGURE 8.6** Decision making under uncertainty. The decision option with the highest expected value will be that which maximizes sigma (*V* x *P*).



and benefits to shutting down a large power generating plant, under the alternative states that either nothing is wrong (and a large expense is incurred in re-starting the plant, and enduring a temporary power loss), or that the plant is failing and will suffer major damage if it continues in operation.

In the analysis of decision making under uncertainty, the exact same procedures as in signal detection theory can be applied for maximizing the expected value of a choice, as long as the probabilities of the different states of the world can be estimated, and as long as values can be placed within the different cells of the matrix (there will be more than four cells if there are more than 2 states or 2 outcomes). The process by which the optimum choice can be proposed involves following calculations analogous to those discussed in the context of Figure 8.5:

1. The probability of each state of the world ($P_S$) is multiplied by the outcome value ($V_{XY}$) in each cell, assigning positive values to "good" outcomes, and negative values to "bad" ones.
2. These [probability X value] products are summed across options, to produce the expected value of each option.
3. The decision alternative with the greatest expected values is chosen.

To the extent that this option is chosen repeatedly over multiple opportunities to exercise the choice, and that values are objective and known, then the algorithm will, *over the long run*, provide the greatest payoff. Such an algorithm, for example, is well suited to apply to a gambling scenario, in which these conditions are met; and it is indeed such an algorithm that is used by gambling casinos to guarantee that they receive a profit (and therefore guarantee that the long term expected value for the gambling consumer is a loss).

While expected value maximization is clear, simple, and objective, there are several factors that complicate the picture when it is applied to most human decisions under uncertainty. First, it is not necessarily the case that people want to maximize their winnings (or minimize their expected losses) over the long run. For example, they may wish to minimize the maximum loss (i.e., avoid picking the option which has a catastrophic negative outcome value). This is, of course, one reason why people purchase fire insurance and avoid the decision option of "no purchase", even though the expected value of the purchase option is negative in the long run (if it were positive for the consumer, the insurance company would soon be out of business!). Second, in many decisions it is not easy to assign objective values like money to the different outcomes. A case in point are decisions regarding safety, in which consequences may be human injury, suffering, or the loss of life. Third, as we discuss in the following section, people do not treat their subjective estimates of costs and values as linearly related to objective values (i.e., of money). Fourth, people's estimates of probability do not always follow the objective probabilities that will establish long term costs and benefits.

In spite of these many departures from the maximum expected value choices in Figure 8.6, departures which we discuss in more detail below, it remains important that we understand the optimal prescription of expected value choices, given that, like the optimal beta, this prescription establishes a benchmark against which the causes of different human departures can be evaluated (Kahneman, 1991), and given the high frequency with which humans make decisions under uncertainty or risk. A few examples are:

- Does the company institute a costly safety program, or does it take gambles that its factory will not be inspected and that an accident will not occur at the workplace?
- Do you purchase the expensive expanded warrantee option for your new computer system: given the likely possibility that it may never fail?
- Does Laura choose the snow over the rock route?

- Does the pilot continue flying through bad weather, or turn back?
- Does the student decide not to read the chapter, gambling that its material will not be covered on the exam?

All of these are examples of risky decision making for which, if probabilities and values are known, the procedures in Figure 8.6 could be applied. We now explore some of the departures or reasons why people make choices that do not agree with the expected value model.

## 6.3 Heuristics and Biases in Uncertain Choice

Whether a choice is between two risky outcomes, or between a risk and a "sure thing" (i.e., an option for which the outcome is known with certainty), decision-making research has revealed a number of ways in which choices depart from the optimum payoff, prescribed by expected value theory. As with diagnosis heuristics, these are not necessarily "bad," and, indeed, some can be shown to be optimal under certain circumstances. Understanding the variables that can moderate the strength of influences on subjective values and probability perception can provide important guidance in improving decision making. We consider below first a shortcut or heuristic related to direct retrieval that totally bypasses the explicit considerations of risk, and then the forms of influences of human perception of value and of probability, which have been incorporated in to a theory of choice known as **prospect theory** (Kahneman & Tversky, 1984).

**6.3.1** DIRECT RETRIEVAL As we have noted in Section 2, many skilled decisions are made without much conscious thought given to risks (probabilities and values). Choices of action may sometimes be implemented simply on the basis of past experience. If the conditions are similar to those confronted in a previous experience, and an action worked in that previous case, it may now be selected in the present case with confidence that it will again produce a satisfactory outcome. This direct retrieval strategy is a hallmark on naturalistic decision making to be discussed below. As well, it is a hallmark of operant conditioning. Indeed studies of decision makers in high stress realistic environments such as fire fighting (Klein, 1997; Klein et al., 1996) reveal the prevalence of such decision making strategies. So long as the domain is familiar to the decision maker, and the diagnosis of the state of the world is clear and unambiguous, the comparative risks of alternatives need not be explicitly considered. Sometimes such an approach may be coupled with a **mental simulation** (Klein & Crandall, 1995), in which the anticipated consequences of the choice are simulated in the mind, to assure that they produce a satisfactory outcome. Good arguments can be made that such a direct retrieval strategy like recognition primed decision making is in fact a highly adaptive one in a familiar domain and if time pressure is high (Svenson & Maule, 1993).

**6.3.2** DISTORTIONS OF VALUES AND COSTS: LOSS AVERSION As we have noted, expected value theory is based upon optimizing some function which in the economic framework has been used to analyze much of human decision making and uses money or objective value as its fundamental currency. But the way that people actually make decisions suggests that they do not view money as a linear function of worth. Instead, much human decision making can be better understood if it is assumed that humans are trying to maximize an expected **utility** rather than expected value (Edwards, 1987), in which utility is the *subjective* value of different expected outcomes. Within this context, the important principle of **loss aversion** specifies that people are more concerned about (greater loss in utility) the loss of a given amount of value, than they appreciate (increase



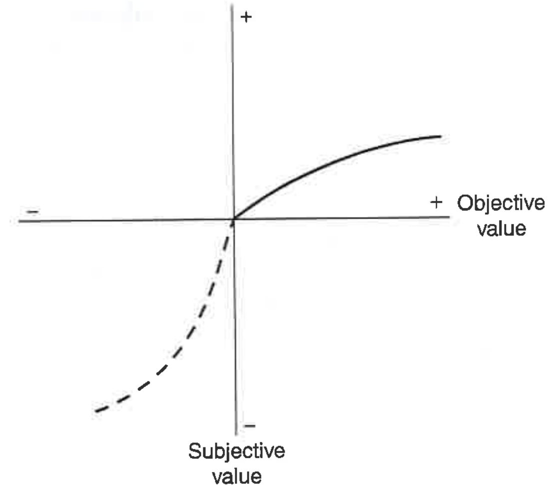


**FIGURE 8.7** The hypothetical relationship between value and utility.

in utility) a gain of the same amount (Garling, 1989; McGraw et al., 2010). This difference is explicitly represented as one important component of the prospect theory of decision making, proposed by Kahneman and Tversky (1984) as shown in Figure 8.7, which relates objective value on the x-axis to subjective utility on the y-axis. To the right, the figure represents the functions for utility gains (receiving money or other valuable items). To the left, it represents the functions for losses. Certain features of this curve nicely account for some general tendencies in human decision making.

The prominent difference in the slope of the positive (gains) and negative (losses) segments of the function represents loss aversion: a potential loss of a given amount is perceived as having greater subjective consequences, and therefore exerts a greater influence over decision-making behavior than does a gain of the same amount. As an example to illustrate this difference, suppose you are given a choice between refusing or accepting a gamble that offers a 50 percent chance to win or lose $1. Most people would typically decline the offer because the potential $1 loss is viewed as more negative than the $1 gain is viewed as positive. As a result, the expected utility of the gamble (as shown in Figure 8.6, the sum of the probability of outcomes times their utilities) is a loss. Another example of loss aversion is what is called the **"endowment effect"** in which people charge more for selling a product (they will lose the product, and their charge is the utility of the loss) than they are willing to pay for it (the utility of the gain, Garling, 1989). The distinct asymmetry between losses and gains appears to reflect operations within different regions of the brain (Lehrer, 2009).

It is important to note that loss aversion is not consistently found, and that the greater impact of losses can sometimes be accounted for by the greater attention paid to and arousal caused by information that anticipates losses (Yechiam & Hoffman, 2012).

A second characteristic of the function in Figure 8.7 is that both positive and negative limbs are curved toward the horizontal as they depart from zero, each showing that equal changes in value produce progressively smaller changes in utility the farther one is from the zero point. This property makes intuitive sense. The gain of $10 if we have nothing at all is more valued than the gain of the same $10 if we already have $100. Similarly, we notice the first $10 we lose, more than an added $10 penalty to a loss that is already $100. Thus, this property captures Weber's Law of Psychophysics as applied to perceived value.

**6.3.3** TEMPORAL DISCOUNTING Differences between value and utility are also reflected in a phenomenon known as **temporal discounting**. Here people often make decisions or chose options that maximize the short term gains (an immediate positive experience) rather than postponing them (a delayed utility) for an option that may result in equal or even greater long term gains; this behavior reflects an implicit belief that the passage of time "discounts" those gains (Mischel Shoda & Rodriguez, 1989). Such behavior seems to explain the attractiveness of borrowing on credit, to obtain an immediate goal (short term gain; Garling, 1989), rather than postponing the goal's receipt until cash is in hand. Temporal discounting appears to differ substantially between people (Ersner-Herschfield et al., 2009). Of course there may be good legitimate reasons to downweight the expected utility of postponing outcomes, in particular because the future is usually uncertain, and less reliably predicted than is the present or immediate future (see discussion of prediction in Chapters 5 and 7). If the probability of future gains is less than of present gains, this difference can offset the greater utility of future gains.

**6.3.4** PERCEPTION OF PROBABILITY We have noted at least three times previously that people's perception of probability is not always accurately calibrated. The "sluggish Beta" phenomenon discussed in Chapter 2, and the representativeness heuristic discussed in this chapter, both illustrated a tendency to downweight the influences of probability in detection and diagnosis, respectively and we introduced the biases in judging proportions in Section 5.1.1. Consistent with these biases in prospect theory, Kahneman and Tversky (1984) have suggested a function relating true (objective) probability to subjective probability (as the latter is inferred to guide risky decision making) that is shown in Figure 8.8.

Four different aspects of this function are critical for understanding risky choice. The first, addressed in Section 5.1.1, is the way in which the probability of rare events are often

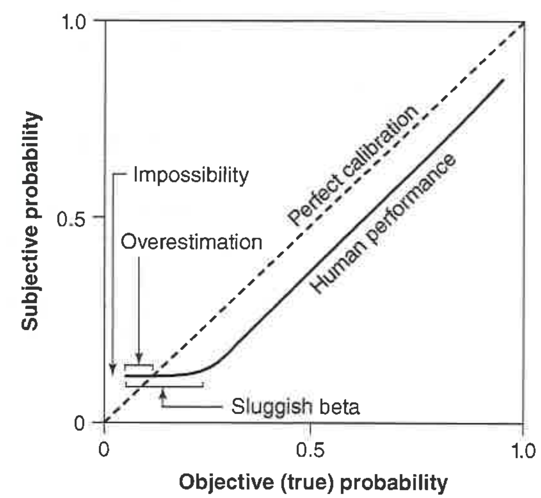


**FIGURE 8.8** A hypothetical weighting function. The solid line represents estimates of subjective probability compared to the perfect calibration of the dashed line.



overestimated, which accounts for two important departures from decision making to maximize expected value: (1) Why do people purchase insurance (choosing a sure loss of money—the cost of the policy—over the risky loss of an accident or disaster, which probably won't happen), and (2) why do people gamble (sacrificing the sure gain of holding onto money for the risky gain of winning)? The answer is that in both cases the risky events are quite rare (the disaster covered by insurance or the winning ticket in the lottery), and hence as shown in Figure 8.8 their probability is subjectively overestimated: The image of winning a gamble looms large, as does the possibility of the disaster for which insurance is purchased. With a larger estimated probability input to the subjective expected utility decision making function, the decision option which anticipates the objectively improbable outcome is more likely to be made.

We do note however, as discussed in Section 5.1.1. that the probability of very rare events may be underestimated if that subjective probability is derived primarily from experience rather than description, and the event in question, because of its rarity, is never personally experienced (Hertwig & Erev, 2009). This second aspect is reflected by the disconnect at the far left of the solid line in Figure 8.8.

The third feature is the relatively lower (than 1.0) slope of the function at its low probability end. This "flat slope" characterizes the reduced sensitivity to probability changes underlying the "sluggish Beta" as well as the representativeness heuristic and ignorance of base rates discussed in Section 5.3.2.

The fourth feature of the function in Figure 8.8 is the fact that for most of its range (i.e., except for the very infrequent events discussed above), the function shows perceived probability as less than actual probability. If the perceived probability that influences one's decision is less than the true probability, then when choosing between two options with positive outcomes, one risky and one certain, the probability of gain associated with the positive risky outcome will be underestimated, and this will also cause the *expected gain* of the risky option to be underestimated; therefore the bias will be to choose the sure thing. When choosing between negative outcomes, the probability of the risky negative outcome will also seem less, the expected **loss** of this option will be **under**-estimated, and it will now be **more** likely to be chosen over the certain loss. It is this third feature, which can be used to account for a very important effect or bias in choice, which is referred to as the **framing effect** or framing bias (Garling, 1989; Kahneman & Tversky, 1984; Mellers, Schwartz, & Cooke, 1998; Munichor, Arev & Lotern, 2006), which we now discuss in detail.

**6.3.5** THE FRAMING EFFECT In its simplest version, the framing effect accounts for how people's preference for outcomes and objects change as function of how their description is *framed* (Tversky & Kahneman, 1981). For example the same ground beef product will seem more attractive if it is described as 80 percent lean than if it is described as 20 percent fat, even though the product is identical in the two descriptions. People will be more likely to choose the beef (over some other meat) with the former description, framed in the positive, than the negative. More seriously, a physician considering treatment of a severely ill patient may have the treatment outcomes listed as a 98 percent chance of survival or a 2 percent chance of mortality. Again, both options describe the same probabilistic outcome. But skilled medical personnel will tend to choose the treatment (over the option, for example doing nothing) more often with the former positive frame than with the 2 percent negative frame (McNeil, Pauker, et al., 1982).

In the above example, we considered the decision to use the treatment (which had a risky, probabilistic outcome) versus doing nothing, whose outcome may be certain. Indeed the framing effect accounts for people's preferences when faced with a choice between a risk and a sure thing. A classic example, faced by most of us at some time or another is when we chose between adhering to some time (or cost) consuming safety procedure (a sure loss), or adopting the risk of avoiding the procedure (driving too fast, running the red light, failing to wear safety glasses) because the **cost of compliance** outweighs our expected benefits of enhanced safety (avoiding the unexpected accident which the safety procedure is designed to prevent). The framing effect as derived from Figure 8.8 accounts for the risk seeking bias when the choice is between the negatives (risk and sure thing), but a risk aversion bias when the choice is between the positives (risk and sure thing; Munichor, Arev, & Lotern, 2006; Simonsohn, 2009).

As a simpler example, if given the choice between winning $1.00 for sure (no risk) and taking a gamble with a 50/50 chance of winning $2.00 or nothing at all (risky)—as we saw above—people typically choose the certain option. They tend to "take the money and run." However, suppose the word "**winning**" was replaced by "**losing**," so that the choice is between losses. This choice produces a so-called avoidance-avoidance conflict, characteristic of the safety decision described above, and people here tend to choose the risky option. They are risk seeking when choosing between losses.

The importance of these differences between perceived losses and gains is that a given change in value (or expected value) may often be viewed *either* as a change in loss or a change in gain, depending on what is considered to be the neutral point or **frame of reference** for the decision making; hence the title of the framing effect. As we saw at the beginning of the chapter, Lauren saw her decision to abandon the summit as a choice between losses. Her teammate gently rephrased this as a choice between gains and this reversed her decision preference. As another example, a tax cut may be perceived as a reduction in loss if the neutral point is "paying no taxes" or as a positive gain if the neutral point is "paying last year's taxes" (Tversky & Kahneman, 1981). As a consequence, different frames of reference used to pose the same decision problem may produce fairly pronounced changes in decision-making behavior (Garling, 1989; Tversky & Kahneman, 1981). Puto, Patton, and King (1985) and Schurr (1987) noted that this kind of bias described the behavior of professional buyers, given hypothetical investment decisions, just as aptly as it described the behavior of typical laboratory subjects. McNeil, Pauker, et al., (1982) found that it also characterized the choices physicians made between safer and riskier treatments.

The effects of framing in an engineering context can be illustrated by considering a process control operator choosing between two courses of action after diagnosing a *potentially* damaging failure in a large industrial process: continue to run while further diagnostic tests are performed or shut down the operation immediately. The first action may be perceived to lead to a very large financial cost (serious damage to the equipment) with some probability much less than 1.0. The second action will produce a substantial cost that is almost certain but of lesser magnitude (start up costs, and lost production time). According to the framing effect, when the choice is framed in this fashion, as the choice between losses, the operator would tend to select the higher-risk alternative (continue to run) over the low-risk alternative (shut down) as long as the expected utilities of the two actions are perceived to be similar. On the other hand, if the operator's perceptions were based on a framework of profits to the company (i.e., gains), the first, risky alternative would be perceived as a probability mix of a full profit if nothing is wrong and a substantially diminished profit if the disastrous event occurs. The second alternative would be perceived as a certain large (but not maximum) profit. Within this positive frame, the choice would be biased toward the second, sure thing alternative: shut the plant down.

The framing effect can also be used to account for the **sunk cost bias** (Arkes & Blumer, 1985; Bazerman, 1998; Molden & Hui, 2011). Here, if we have made a bad decision, perhaps a poor investment, and have already lost a great deal, then when confronted with the choice of whether to "get out" and cut the losses, rather than continue with the investment, people will be more likely to continue ("throw good money after bad"), even when it is in their economic interest to withdraw (a lower expected loss). Rationally, the previous history of investment should not enter into the decision for the future. Yet it does. People faced with the exact same choice but when they were **not** responsible for the initial investment decision (that had lost utility) will be far more inclined to cut their losses and choose to terminate the investment (a sure loss). We can see how this was illustrated by Lauren's initial decision to push on toward the summit.



The interpretation of the sunk cost bias within the framing context is straightforward. For the investor whose previous decision was poor, the choice is between a sure loss (get out now) and a risky loss (the bad investment *may* turn good in the future, but is more likely to continue to worsen). For the newcomer, encountering the same situation, but whose own utility had not been diminished by the bad decision, the "sure thing" option is neither loss nor gain. Hence the choice is between 0 utility and an expected loss; a circumstance that fairly easily leads to a bias to choose to terminate the investment.

### 6.4 The Decision to Behave Safely

The phenomenon of framing applies to a wide variety of risky choices made by people in society. As we have noted, a common choice is whether or not to adhere to a particular safety regulation; wearing a seatbelt, a protective helmet or harness, or some other behavior in the workplace. The sure "cost of compliance" is always explicitly or implicitly compared against the expected negative utility of the more risky behavior. In making such choices, it is important to bear in mind the influence of the framing effect—to the extent that outcomes are viewed as negatives, the risky behavior may be chosen more often—as well as the two related heuristics discussed in Section 5.3 which influence diagnosing the state of the world:

The availability heuristic indicates that the perceived frequency of different negative consequences of unsafe behavior will be based not on their actual frequency (objective risk), but upon their salience in memory, if those consequences were either directly experienced or learned through description. When these do *not* correspond, risks can be seriously misestimated. The representativeness heuristic (and base rate ignoring) suggests that we may not be very sensitive to the probability of disastrous consequences at all; and indeed a study by Young, Wogalter, and Brelsford (1992) found that the perceived **severity** of a hazard has a greater impact on risk estimation than does the probability of the hazard. Finally, it is the case that both perceived severity and probability will be **abstract** experiences in making the choice, only possibly perceived in the future. As temporal discounting suggest (see 6.3.3), their expected costs may be diminished. In contrast, the cost of compliance imposes a direct tangible and present experience (e.g., the discomfort of wearing a safety device or the inconvenience of adhering to safety procedure), the experience is highly **accessible** (Kahneman & Frederick, 2002). This analysis suggests that risk mitigation efforts should be directed heavily to reducing the cost of compliance more than increasing the perceived negative risks of the accident.

In addition to the "sure-thing versus risk" choice to behave unsafely, people also allow risks to enter into their everyday safety decisions by balancing perceived risks, for example in their choice of transportation modes, in foods to eat, or to behave in a way that is sensitive to climate change (Dotta, 2011). In analyzing such behavior, it is important to realize the substantial departures between people's perception of relative risks and the true measures of risk (as for example defined by probability of death). As one example, the probability of death from a fall in the home, is far more likely than the probability of death from an airplane crash; but people's perception of these risks are often reversed (Combs & Slovic, 1979).

At least three factors appear to be responsible for the fact that people elevate their estimate of risk above the true "objective" values associated with, for example probability of death. The first is the fact that publicity, for example from the news media, tends to make certain risks more

**available** to memory than others (Combs and Slovic, 1979). Hence we observe the high perceived risks of well publicized events (like a major plane crash or a terrorist bombing). Second, people's perception of risk is driven upward by what is described as a "dread factor" (uncontrollable, catastrophic consequences, unavoidable), and third, perceived risk is inflated by an "unknown" factor, which characterizes the risk of new technology, such as genetic manipulations and many aspects of automation (Slovic, 1987).

It is important for policy makers to consider these influences on the risk perceived by the public. But it is equally important for all people who make choices based upon risk to consider the consequences of those choices on scarce resource allocation (Keeney, 1988). For example the choice to allocate a large amount of money to reduce one particular risk, whose objective risk is small (but perceived riskiness is large), may be made at a cost of pulling those resources away from mitigating a much larger objective risk, whose subjective perception is smaller.

An important way to mitigate risky behavior when it results because the probability of the negative event may be very rare (and hence never personally experienced) even as its negative consequences may be severe, is through "gentle reminders" (Hertwig & Erev, 2009). This technique imposes minor penalties—a gentle reminder—for the risk-producing *behavior* (e.g., failing to heed a safety precaution) which will be *experienced* much more frequently than the rare severe consequences. Such a technique has proven effective in inducing more safety compliant behavior in hospitals.

In conclusion, we note that risk perception, and risk seeking are influenced by a host of other factors, besides the framing of negative outcomes. For example time stress appears to lead to more risk seeking (Chandler & Ponin, 2012), and Figner and Weber (2011) discuss other contextual and individual difference factors that influence risk seeking.

## 7. EFFORT AND META COGNITION

Our treatment of decision making up to now has focused most on the extermal drivers of decision making—problem structure, risk, values, and probability—as filtered by human cognition. However, as shown in Figure 8.1, there are two critical inputs to the decision process emanating from the decision maker himself or herself: effort and meta-cognition. Because these two are interrelated, we treat them together as follows, even as meta-cognition was discussed in the previous chapter, and effort in Chapter 10.

### 7.1 Effort

In our discussion of decision fatigue, we emphasized that effective decision making often requires effort. Resource-dependent working memory is necessary to diagnose and evaluate options. Decision making competes for those resources with concurrent tasks (e.g., Sarno & Wickens, 1995; see Chapter 10), and sustained decision making depletes that pool of resources or cognitive effort (Tierney, 2011). Indeed, it has been shown that repeated decision making competes with the effort required for exerting self control in other aspects of life (e.g., resisting temptations; Tierney, 2011). Not surprisingly then, the variety of decision-making strategies will vary in their effort requirements (Bettman, Johnson, & Payne, 1990; Johnson & Payne, 1985; Payne, Bettman, & Johnson, 1993). In particular, heuristics, such as elimination-by-aspects or representativeness, can be viewed as "effort-lite" versions of the more accurate, full compensatory choice model (Section 6.1) or base rate consideration (Section 5.4.1), respectively.

The effort required and accuracy observed of these two classes of DM strategies is reflected schematically in Figure 8.9, which indeed previews the concept of the **performance-resource function**, to be discussed in Chapter 10. Within this context, effort itself can be viewed as a



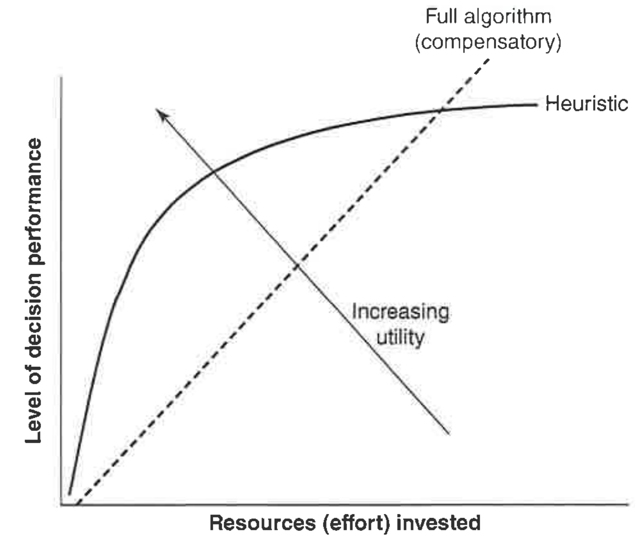


**FIGURE 8.9** Effort, performance and heuristics in decision making. The figure shows the improvement in decision performance as a function of more effort invested into the decision process for heuristics (solid line) and algorithms (dashed line). With small effort investment, heuristics can produce better performance than algorithms.

valuable resource to be conserved. For example, as more resources are invested, performance with both elimination-by-aspects heuristic and the compensatory algorithm will improve. However with a small investment of resources, the "efficiency" of decision making (accuracy per resources invested) will be greater with the heuristic; and greater efficiency can be considered as more optimal, when time or resources are scarce. Time pressure will place greater premium on effort conservation. Thus the pilot who dithers in deciding what to do, as the plane heads toward a hillside or is running out of fuel, will surely be considered to be non-optimal (Orasanu & Fischer, 1997). The **contingent model** of decision strategies developed by Payne, Bettman, and Johnson, 1993, predicts how different strategies will be chosen, contingent upon the available time (resources).

Another important example of this *contingency* of decision strategy choice upon effort and accuracy requirements is in the choice of whether to terminate a diagnosis or seek further (often confirmatory) evidence, given the effort required for further information access (see also search termination discussed in Chapter 3 Section 2.1). For example, in deciding whether or not a particular set of findings warrant inclusion as a general principle in this text, the authors make decisions on whether it is worth the effort and time to go back to the library and do further information search regarding the findings in question. What will be the perceived gain in seeking more information (MacGregor, Fischoff, & Blackshaw, 1987)? How much time will it take me to do so? How confident am I now that I have made an appropriate diagnosis of the state of human performance already, to include the principle in question as part of a chapter?

Of course the tradeoff between accuracy and effort in choosing a strategy is not always based on the actual level of these variables, but instead is based on the *anticipated* accuracy and effort (Fennema & Kleinmuntz, 1995; Seagull Xiao & Plasters, 2004). In this regard, research has revealed that people are not fully calibrated, in relating the anticipation of accuracy and effort, to the *actual* accuracy achieved and effort experienced (Fennema & Kleinmuntz, 1995).

### 7.2 Meta-Cognition and (Over) confidence

The issues of anticipated effort and accuracy, and the conscious choice of a decision strategy brings us to the important role of meta-cognition in decision making. What does the decision maker know (or think) about the accuracy of his diagnosis and choice? How does this anticipation influence the choice of strategy and subsequent decision-making behavior (including the choice not to decide at all, as in the case of the parole boards discussed in Section 5.4.3). As Kahneman and Klein (2009) note, this is the role of the type 2 system: to oversee, review, and audit the more automatic decision-making behavior of the type 1 system.

It turns out that one of the most critical and enduring influences on meta-cognition is the confidence in assessing ones own diagnosis and judgment. Such confidence is often unrealistically high, as manifest in the overconfidence bias (Nickerson, 1998). In diagnosis, confidence judgments will influence the extent to which we jump into action (choice), rather than seek more evidence, or prepare for the case in which the assessment may have been wrong. In choice, confidence assessments will influence the extent to which we plan for alternative actions (to the extent that we think our chosen action might have been wrong). In both cases, as Griffin and Tversky (1992) state: "although overconfidence is not universal, it is prevalent, often massive, and difficult to eliminate". Several examples from different walks of life may be cited:

- The average driver estimates him/herself to be within the top 25th percentile of safe drivers (Brehmer, 1981). By definition, if confidence were calibrated, this should be 50 percent.
- Fischoff (1977) and Fischhoff and MacGregor (1982) asked people to make prediction about future events (e.g., elections, winners of athletic contests), and noted that, typically, whereas predictions might turn out to be 60 percent accurate (evaluated after the event took place), the confidence offered as to prediction accuracy would be more like 80 percent.
- Such overconfidence is not confined to novices in a field, as Tetlock (2005) performed a long term study of experts in political forecasting, and observed similar overconfidence. This was just as prevalent and severe as in novices making similar predictions.
- OC is well documented in the **planning fallacy** (Buehler Griffin & Ross, 2002). Here people are eternally optimistic in their projections of how long it will take (or how many resources will be required) to do something, from achieving a personal goal (like turning in an assignment on time), to completing massive construction projects like the Denver International Airport or the Sidney Opera house. Indeed in one study, students expressed 84 percent confidence that they would complete an assignment on time, whereas in fact, only 40 percent did so (Buehler et al, 2002).
- Scientists are notoriously overconfident about the precision of their estimates of various physical constants, such as the speed of light (Henrion & Fischoff, 2002).
- Sulistyawati, Wickens, and Chui (2011) observed that those pilots who showed more overconfidence in their situation awareness estimates were in fact less accurate in those estimates.

We have also encountered OC in other chapters: in Chapter 3, this was illustrated by the phenomenon of "change blindness blindness" (Levin, Momen, et al., 2000), which describes people's overconfidence in their ability to detect unexpected events. In Chapter 5, we considered their overconfidence in detecting hazards at night, leading to overspeeding (Leibowitz, Post et al., 1982). In eye witness testimony, discussed in Chapters 2 and 7, we learned of the general tendency to be overconfident of the accuracy of their own recognition memory (Brewer & Wells, 2006; Wells, Lindsay, & Ferguson, 1979), and in learning itself (Chapter 7) people tend to allow the ease of learning to act as a proxy for the ease of later recall (it is not), and hence be overconfident in the accuracy of their predicted level of recall (how well they will do on the test), thereby underestimating



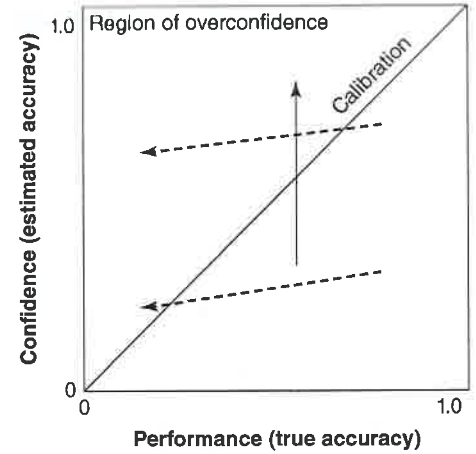


**FIGURE 8.10** How confidence and overconfidence is driven by reliability. Each arrow represents the effect of some task variable on both decision accuracy and performance, as described in the text.

their need for study (Bjork, 1999). In Chapter 10 we will encounter OC again in the context of people's confidence in their ability to time share while driving (Horrey, Lesch, & Garabet, 2009).

Of course there is great variability between individuals and circumstances in the extent to which OC is manifest, and we describe below some key moderating variables. First however we can formally represent OC within the context of the **accuracy-confidence calibration space** as shown in Figure 8.10. When confidence is expressed by predicted or judged accuracy (e.g., how well do you think you did on the test), then the two variables, actual and predicted performance, can be presented on the same scale, and this graph can define the region of OC as shown above and to the left of the diagonal line of perfect calibration. Furthermore Figure 8.10 illustrates a relatively common phenomenon by the dashed arrows, in which a variable that diminishes accuracy *fails* to produce a parallel loss of confidence, and we see that this phenomenon often (top dashed arrow) but not invariably (bottom dashed arrow) leads to OC. Somewhat less prevalent is the pattern represented by the solid arrow, in which a variable influences confidence, even as accuracy is little changed.

Research has now identified several variables that create overconfidence including the following in particular:

1. Diagnostic or problem difficulty. This effect can be described in different ways. For example when two hypotheses become less discriminable (more ambiguous cues) accuracy of diagnosis decreases, but confidence does not, echoing the pattern shown in the upper dashed arrow in Figure 8.10 (Fishcoff, 1977; Koehler, Brenner, & Griffin, 2002). Evaluating pilots' diagnosis of aviation problems, Mosier et al. (2007) found a relation paralleling that line. In domains where prediction is hard to make accurately because of many uncertainties (stock brokers, politics, mental health), overconfidence is prevalent (even by experts) whereas it is less so in more predictable domains like weather forecasting (Kahneman & Klein, 2009; Taleb, 2007; Tetlock, 2005). So too, poorer drivers (for which driving is, by definition, a more difficult task) show more OC than better drivers (Kidd & Monk, 2009).
2. Evidence reliability. People are not very sensitive to differences in evidence reliability (as we saw with the as-if heuristic; Griffin & Tversky, 1992) and are guided more by the strength of evidence than by its reliability. Thus when reliability and performance decline (e.g., by samples with smaller N), their confidence in the impact of the message provided by this

lower reliability (lower information value) does not. These changes all reflect differences along the upper dashed arrow.

3. In a pattern reflecting the solid arrow of Figure 8.10, when people rely on progressively more sources of *correlated* information, they gain confidence (Kahneman & Klein, 2009). The problem is, when information is highly correlated, errors (unreliability) in one source will typically co-occur in other sources (e.g., a common failure may underlie both), and so confidence should not proportionately increase. For example, consider two witnesses both depending on the same, unreliable source of hearsay evidence.
4. Progressively more sources of information (whether correlated or not) will typically increase confidence in a diagnosis. But as we discussed in Section 5.2.2, this often does not lead to an increase in diagnostic accuracy.

In the above discussion of OC, we have examined differences in conditions that may differentially influence confidence and accuracy. But we can also ask about differences between people. Are there certain classes of people whose performance tends to occupy the upper left portion of the space. This issue is of particular relevance to assessments of the accuracy of judicial eye witness testimony (Hope et al., 2004).

## 8. EXPERIENCE AND EXPERTISE IN DECISION MAKING

As we discussed earlier in this chapter, experts often (but not always) make better decisions than novices. As we have noted above, this phenomenon is well captured by the study of naturalistic decision making (Kahneman & Klein, 2009; Mosier & Fischer, 2010; Montgomery, Lipshitz, & Brenner, 2005; Zsambok & Klein, 1997), which captures the experience-related differences associated with the two major stages of decision making. In front end decision making (diagnosis), experts typically manifest recognition primed decision making (RPDM). Here through repeated exposure to the same set of *correlated* cues, leading to the same state assessment, experts are able to automatically classify the appropriate state, almost the same as the automatic pattern recognition discussed in Chapter 6. Hammond et al. (1987) refer to this as **holistic decision making**, a function associated with decision system 1 (Kahneman & Klein, 2009). Schriver, Morrow, Wickens, and Talleur (2008) found that expert pilots were better able to exploit correlated cues in airplane fault diagnosis than were novices. Their decision advantage was less pronounced when cues were uncorrelated.

Also, as noted in Section 6.3.1, in back end decision making, experts can accomplish direct retrieval of choices from long-term memory quite rapidly. What often worked before (given a RPDS situation assessment; a good outcome) will work again. This phenomenon was observed by the more rapid response shown by expert pilots by Schriver et al. (2009) with no sacrifice of accuracy. And yet, as we have seen, the success of expertise in DM is far from guaranteed (Tetlock, 2005). Cues may be uncorrelated, overconfidence may short change meta cognitive monitoring, and rapid pattern- recognition classification may overlook a single outlying cue.

Furthermore, as we have considered before, practice in decision making does not necessarily make perfect, as it does in other skills. Expertise in some decision-making tasks does not guarantee immunity to certain biases and heuristics (Kahneman & Klein, 2009; Taleb, 2007; Tetlock, 2005). Some assistance in solving the puzzle as to why experienced decision makers are neither perfect nor sometimes better than novices is provided by Shanteau's (1992) careful classification of the domains and properties of those domains, that distinguish when expertise *does* develop from practice, and when it does not (Table 8.1).

Kahneman and Klein (2009) in particular, have highlighted the extent to which expertise in decision making (where experience helps) only emerges in domains such as weather forecasting,



**TABLE 8.1 From Shanteau (1992)**

| | |
|---|---|
| Domains of "Good" Decision Making | Domains of "Poor" Decision Making |
| Weather Forecasting | Clinical Psychologists |
| Chess Masters | Personnel Selectors |
| Physicians | Parole Officers |
| Photointelligence Analysts | Stock Brokers |
| Accountants | Court Judges |
| Characteristics of the Domains: | |
| Dynamic | Static |
| Decisions About Things | Decisions About People |
| Repetitive | Less Predictable |
| Feedback Available | Less Feedback |
| Decomposable Decision Problems | Not Decomposable |

in which the pattern of cue correlations is relatively strong, and different predicted states can be well discriminated.

So why does decision making not improve much with experience in these other cases? Einhorn and Hogarth (1978) have added insight to understanding the problems of learning in decision making; characteristic of the right side of Table 8.1; by addressing the role of *feedback* in the typical decision-making problem. As we noted in Chapter 7, feedback is critical for nearly any form of learning or skill acquisition. Yet several characteristics of decision making prevent it from offering its usual assistance.

1. **Feedback is often ambiguous**, in a probabilistic or uncertain world. That is, sometimes a decision process will be poorly executed, but because of good luck will produce a positive outcome; at other times, a decision process can follow all of the best procedures, but bad luck produces a negative outcome. In the first case, the positive reinforcement will increase reliance on the bad process, whereas in the second case, the punishment, realized by the bad outcome, will extinguish the effective processing that went into the decision.
2. **Feedback is often delayed**. In many decisions, such as those made in investment, or even prescribing treatment in medicine, the outcome may not be realized for some time. As we discussed in Chapter 7, added delay in feedback beyond a few minutes is rarely of benefit. In decision making the reason is that, when the feedback finally arrives, the decision maker may have forgotten the processes and strategies used to make the decision in the first place, and therefore may fail to either reinforce those processes (if the feedback was good) or correct them (if the feedback was bad). Furthermore, because feedback is delayed, decision makers may well have turned their attention to other problems and provide less attention to processing it than they would if feedback arrived immediately after. Finally, in a phenomenon that we know as "Monday morning quarterbacking" or "hindsight bias," Fischhoff (1977) and Woods et al. (1994) have documented the extent to which, after an outcome is known, we revise our memory of what we knew before the decision was made in such a way as to downplay our "surprise" at its outcome ("I knew it all along"). If we do not consider ourselves surprised by the outcome (in hindsight), then we will foresee less reason to revise our decision process (i.e., learn from the outcome).
3. **Feedback is processed selectively.** Einhorn and Hogarth have considered the learning of a decision maker who is classifying applicants as either acceptable to or rejected from a

program, and is learning from feedback, regarding the outcome of those who were selected (see Figure 8.11). As the decision maker may process feedback from this process, we note that he or she will typically only have available feedback from those who were admitted (and succeeded or failed), rarely learning if the people excluded by his decision-making rule would have succeeded had they been admitted. Furthermore, the confirmation bias will tend to lead people to focus more attention on those who were admitted and succeeded (therefore confirming that the decision rule was correct), than those who were admitted and failed (therefore disconfirming the validity of the decision rule). As shown in Figure 8.11b, they may provide extra assistance to those admitted by their rule—influencing the outcome of the decision to provide further confirmation of the correctness of the rule.

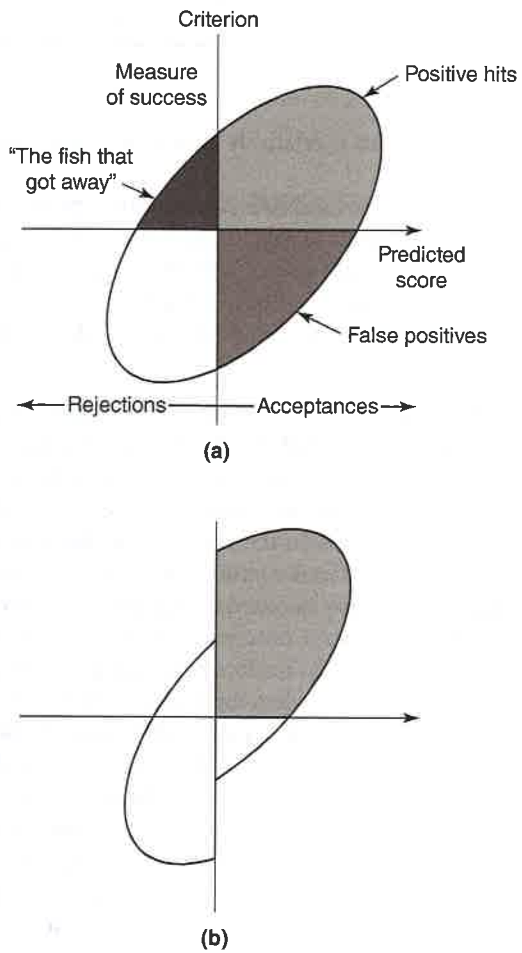


**FIGURE 8.11** (a) Source of unwarranted confidence in prediction. A predicted score of applicants, reflecting the decision maker's rule, is shown on the x-axis. The actual measure of success is shown on the y-axis. (b) The influence of extra assistance to those admitted to the program. *Source:* H. J. Einhorn and R. M. Hogarth, "Confidence in Judgment: Persistence of the Illusion of Validity," *Psychological Review, 85* (1978), p. 397. Copyright 1978 by the American Psychological Association. Adapted by permission of the authors.



## 9. IMPROVING DECISION MAKING

In reviewing the material we have covered in this chapter, one may characterize human decision making as either generally "good" (by focusing on its many successes) or "faulty" (by focusing on its failures). While we have no interest in taking a stand on this scale of evaluating human decision making; we believe that as long as there is evidence that some decision making can be improved in some circumstances, it is the responsibility of engineering psychology to recommend possible ways of supporting that improvement. We consider four such techniques in this chapter related to training, proceduralization, displays, and automation.

### 9.1 Training Debiasing

As we saw above, pure practice at decision making is not necessarily an effective or efficient way of improving its quality. Instead, research has focused on more targeted practice and instructions to remove or reduce many of the biases discussed above, a technique known as **debiasing** (Fischhoff, 1977; 2002; Larrick, 2006; Lipshitz & Cohen, 2005). In a review of debiasing literature, Larrick (2006) concluded that pure instructions or exhortations to avoid biases are ineffective. Correspondingly he found little evidence that simply teaching people *about* biases (e.g., reading this chapter) is effective. This may represent "inert knowledge" which can be undertstood, but not transferred to practice. Instead, effective techniques focus not only on instructing the nature of a particular bias in question, but providing specific examples, and practicing the debiasing strategies (Fong et al., 1991). The following are some specific examples of success.

Hunt and Rouse (1981) have succeeded in training operators to extract diagnostic information from the absence of cues. In sequential cue information integration tasks, Lopes (1982) and Wickens, Ketels, et al. (2010) successfully reduced anchoring through training, the latter instructing subjects about the reduced reliability of older information (see 5.5.1).

Some success in reducing the confirmation bias has also been observed by the training strategy of "consider the opposite" (Mussweiler et al., 2000). For example Koriat, Lichtenstein, and Fischhoff (1980), and Cohen, Freeman, and Thompson (1997) have both found that forcing forecasters to entertain reasons why their forecasts might **not** be correct reduced their biases toward overconfidence in the accuracy of the forecast.

Also successful is a kind of training aid designed to provide more comprehensive and immediate feedback in predictive and diagnostic tasks, so that operators are forced to attend to the degree of success or failure of their rules. We noted that the feedback given to weather forecasters is successful in reducing the tendency for overconfidence in forecasting (Murphy & Winkler, 1984). Jenkins and Ward (1965) demonstrated that providing decision makers simultaneously with data in all four outcomes of a decision like that represented in Figure 8.11 instead of simply the hit probability, improves their appreciation of predictive relations. Where selection tasks or diagnostic treatments are prescribed, box scores should be maintained to integrate data in as many cells of the matrix as possible (Einhorn & Hogarth, 1978; Goldberg, 1968). Tversky and Kahneman (1974) suggested that decision makers should be taught to encode events in terms of probability rather than frequency since probabilities intrinsically account for events that did not occur (negative evidence) as well as those that did.

Finally, in an interesting take on debiasing training, Fischhoff (2002) described the success of some training programs designed to reduce the prevalence of teens engaging in risky behavior (drinking, speeding). Here he makes the point that such behavior, while actually not very frequent, is highly salient, compared to the prevalence of safe behavior. As we have noted above, salient but rare described events may be overestimated in their frequency. If training programs emphasize instead the high frequency of teens engaged in safe behavior, the peer-pressure tendency to imitate the latter (e.g., behave safely) is increased.

### 9.2 Proceduralization

While debiasing is a form of training that often focuses people's awareness directly on understanding the sources of their cognitive limitations, **proceduralization** simply outlines prescriptions of techniques that should be followed to improve the quality of decision making (Bazerman, 1998). This may include for example prescriptions of following the decision decomposition steps of diagnosis and choice theory, as shown in Figures 8.5 and 8.6 (Larrick, 2006). Such a technique has been employed successfully in certain real world decisions which are easily decomposable into attributes and values, such as selecting the location of the Mexico City airport (Kenney, 1973), or assisting land developers and environmentalists to reach a compromise on coastal development policy (Gardner & Edwards, 1975). The formal representation of fault tree and failure modes analysis (Kirwan & Ainsworth, 1992; Wickens, Lee, et al., 2004), is a procedure that can assist the decision maker in diagnosing the possibility of different kinds of system failures. A study of auditors by Ricchiute (1998) has recommended a procedure by which evidence, accumulated by a junior auditor, is compiled and presented to a senior auditor who makes decisions, in such a way as to avoid the sequential biases often encountered in processing information (see Section 5.4).

In a way that integrates debiasing training and proceduralization, Leher (2010) has summarized research to suggest five strategies for effective decision making:

1. Simple problems require reasoning. (Using ones "gut" reflected in the type 1 system, may be a part of this, but type 2 system analysis can almost always help).
2. Novel problems require reasoning. Given the type 1 system may not be available here, it is important to examine past experience analytically to determine how these past decisions might advise a current, complex decision.
3. Embrace uncertainty. Always entertain competing hypotheses. Always remind yourself of what you *don't* know.
4. You know more than you know. *Once* you have developed some level of expertise in an area, *then* it is OK to trust your emotions and your "gut", which can reflect the massively parallel processes in the brain to suggest that certain choices may "seem right" and others are troublesome. But the type 2 system needs to audit these gut calls.
5. Think about thinking: the advocacy of meta-cognition.

### 9.3 Displays

There is good evidence that effective displays can support the front end of decision processes (cue integration and diagnosis), by assisting the deployment of selective attention (Mosier & Fischer, 2010). For example, Stone, Yates, and Parker (1997) observed that pictorial representations of risk data supported more calibrated risk decisions than do numerical or verbal statements. Schkade and Kleinmuntz (1994), studying the decision processes of loan officers, found that the format in which information regarding the attributes of different loan applicants was structured influenced the nature of the judgments in a way suggesting that people minimized the amount of attentional effort required for information integration. Cook and Smallman (2008) found that an integrated graphical display of intelligence cues shown to professional intelligence analysis reduced the confirmation bias, relative to a text-based presentation which implicitly suggested a sequential ordering (and hence invited sequential biases).



The proximity compatibility principle (Wickens & Carswell, 1995), described in Chapter 3, is relevant to effective decision making, prescribing that sources of information that need to be integrated in diagnosis are made available simultaneously (not sequentially) and in close display proximity to each other so that all can be accessed with minimal effort. Emergent features of object displays can sometimes facilitate the integration process in diagnosis (Barnett & Wickens, 1988). In this regard, we also saw in Chapter 4 that ecological displays assisted professionals in the diagnosis stage of process control fault management, corresponding to front end decision making (Burns et al., 2008).

### 9.4 Automation and Decision Support Tools

Finally, automation and expert systems have offered promise in supporting human decision making. This is described in much more detail in Chapter 12, but to provide a link here, such support can be roughly categorized into front end (diagnosis and situation assessment) and back end (treatment, choice, and course-of-action recommendations) support. This dichotomy is well illustrated in the two major classes of medical decision aids (Garg et al., 2005; Morrow, Wickens, & North, 2005), where both have enjoyed some modest success. We also note here that procedures whereby humans estimate weights and cue values for diagnostic problems, but computers perform the integration of those values (e.g., Dawes & Corrigan, 1974; Fischhoff, 2002) dictate a preferred allocation of function between human and automation in a cooperative human-automation decision endeavor.

## 10. CONCLUSION AND TRANSITION

In conclusion, we see that decision making is complex and interactive, with different components invoking common cognitive and information processing mechanisms (e.g., overconfidence in both diagnosis and choice). The topic also links to earlier topics of attention, perception, and memory, as well as the topic of limited resources that we will discuss in Chapter 10. At this time it is appropriate to turn our attention to decisions of a more rapid and automatic sort, often studied in the laboratory in the context of reaction time. Thus our focus now in Chapter 9 will be on the decisions that select and execute rapid actions, under some degree of time pressure.

## Key Terms

absence of a cue *256*
accessibility *259*
accuracy-confidence calibration space *277*
anchoring heuristic *260*
as-if heuristic *256*
attribute substitution *260*
availability heuristic *259*
base rate *259*
Bayesian *260*
choice *248*
choice of action *262*
confirmation bias *261*
cost of compliance *272*
Cue diagnosticity *254*
Cue reliability *254*
debiasing *281*
decision fatigue *263*
diagnosis *248*
elimination-by-aspects *265*
endowment effect *269*
expected value *249*
extrapolating non-linear trends *251*
frame of reference *272*
framing effect *271*
gambler's fallacy *252*
heuristics/biases *246*
hindsight bias *250*
holistic decision making *278*
information processing *247*
information value *254*
loss aversion *268*
mental simulation *268*
meta-cognition *249*
naturalistic decision making *246*
normative decision making *247*
overconfidence bias *263*
performance-resource function *274*
planning fallacy *276*
prevalence rates *259*
primacy *261*
proceduralization *282*
prospect theory *268*
representativeness heuristic *258*
risk *246*
salience bias *256*
satisficing *265*
sunk cost bias *272*
temporal discounting *270*
uncertainty *246*
utility *268*