
A Perspective on Judgment and Choice

Mapping Bounded Rationality

Daniel Kahneman
Princeton University

Early studies of intuitive judgment and decision making conducted with the late Amos Tversky are reviewed in the context of two related concepts: an analysis of accessibility, the ease with which thoughts come to mind; a distinction between effortless intuition and deliberate reasoning. Intuitive thoughts, like percepts, are highly accessible. Determinants and consequences of accessibility help explain the central results of prospect theory, framing effects, the heuristic process of attribute substitution, and the characteristic biases that result from the substitution of nonextensional for extensional attributes. Variations in the accessibility of rules explain the occasional corrections of intuitive judgments. The study of biases is compatible with a view of intuitive thinking and decision making as generally skilled and successful.

The work cited by the Nobel committee was done jointly with the late Amos Tversky (1937–1996) during a long and unusually close collaboration. Together, we explored a territory that Herbert A. Simon had defined and named—the psychology of bounded rationality (Simon, 1955, 1979). This article presents a current perspective on the three major topics of our joint work: heuristics of judgment, risky choice, and framing effects. In all three domains, we studied intuitions—thoughts and preferences that come to mind quickly and without much reflection. I review the older research and some recent developments in light of two ideas that have become central to social-cognitive psychology in the intervening decades: the notion that thoughts differ in accessibility—some come to mind much more easily than others—and the distinction between intuitive and deliberate thought processes.

The first section, Intuition and Accessibility, distinguishes two generic modes of cognitive function: an intuitive mode in which judgments and decisions are made automatically and rapidly and a controlled mode, which is deliberate and slower. The section goes on to describe the factors that determine the relative accessibility of different judgments and responses. The second section, Framing Effects, explains framing effects in terms of differential salience and accessibility. The third section, Changes or States: Prospect Theory, relates prospect theory to the general proposition that changes and differences are more accessible than absolute values. The fourth section, Attribute Substitution: A Model of Judgment by Heuristic, pre-

sents an attribute substitution model of heuristic judgment. The fifth section, Prototype Heuristics, describes that particular family of heuristics. A concluding section follows.

Intuition and Accessibility

From its earliest days, the research that Tversky and I conducted was guided by the idea that intuitive judgments occupy a position—perhaps corresponding to evolutionary history—between the automatic operations of perception and the deliberate operations of reasoning. Our first joint article examined systematic errors in the casual statistical judgments of statistically sophisticated researchers (Tversky & Kahneman, 1971). Remarkably, the intuitive judgments of these experts did not conform to statistical principles with which they were thoroughly familiar. In particular, their intuitive statistical inferences and their estimates of statistical power showed a striking lack of sensitivity to the effects of sample size. We were impressed by the persistence of discrepancies between statistical intuition and statistical knowledge, which we observed both in ourselves and in our colleagues. We were also impressed by the fact that significant research decisions, such as the choice of sample size for an experiment, are routinely guided by the flawed intuitions of people who know better.

Editor's note. This article is based on the author's Nobel Prize lecture, which was delivered at Stockholm University on December 8, 2002, and on the text and images to be published in *Les Prix Nobel 2002* (Frängsmyr, in press). A version of this article is slated to appear in the December 2003 issue of the *American Economic Review*.

Author's note. This article revisits problems that Amos Tversky and I studied together many years ago and continued to discuss in a conversation that spanned several decades. The article is based on the Nobel lecture, which my daughter Lenore Shoham helped put together. It draws extensively on an analysis of judgment heuristics that was developed in collaboration with Shane Frederick (Kahneman & Frederick, 2002). Shane Frederick, David Krantz, and Daniel Reisberg went well beyond the call of friendly duty in helping with this effort. Craig Fox, Peter McGraw, Daniel Oppenheimer, Daniel Read, David Schkade, Richard Thaler, and my wife, Anne Treisman, offered many insightful comments and suggestions. Kurt Schoppe provided valuable assistance, and Geoffrey Goodwin and Amir Goren helped with scholarly fact-checking. My research is supported by National Science Foundation Grant 285-6086 and by the Woodrow Wilson School for Public and International Affairs at Princeton University.

Correspondence concerning this article should be addressed to Daniel Kahneman, Department of Psychology, Princeton University, Princeton, NJ 08544-1010. E-mail: kahneman@princeton.edu

Daniel Kahneman



In the terminology that became accepted much later, we held a two-system view, which distinguished intuition from reasoning. Our research focused on errors of intuition, which we studied both for their intrinsic interest and for their value as diagnostic indicators of cognitive mechanisms.

The Two-System View

The distinction between intuition and reasoning has been a topic of considerable interest in the intervening decades

(among many others, see Epstein, 1994; Hammond, 2000; Jacoby, 1991, 1996; and numerous models collected by Chaiken & Trope, 1999; for comprehensive reviews of intuition, see Hogarth, 2001; Myers, 2002). In particular, the differences between the two modes of thought have been invoked in attempts to organize seemingly contradictory results in studies of judgment under uncertainty (Kahneman & Frederick, 2002; Sloman, 1996, 2002; Stanovich, 1999; Stanovich & West, 2002). There is considerable agreement on the characteristics that distinguish the two types of cognitive processes, which Stanovich and West (2000) labeled *System 1* and *System 2*. The scheme shown in Figure 1 summarizes these characteristics: The operations of System 1 are typically fast, automatic, effortless, associative, implicit (not available to introspection), and often emotionally charged; they are also governed by habit and are therefore difficult to control or modify. The operations of System 2 are slower, serial, effortful, more likely to be consciously monitored and deliberately controlled; they are also relatively flexible and potentially rule governed. The effect of concurrent cognitive tasks provides the most useful indication of whether a given mental process belongs to System 1 or System 2. Because the overall capacity for mental effort is limited, effortful processes tend to disrupt each other, whereas effortless processes neither cause nor suffer much interference when combined with other tasks (Kahneman, 1973; Pashler, 1998).

As indicated in Figure 1, the operating characteristics of System 1 are similar to the features of perceptual processes. On the other hand, as Figure 1 also shows, the operations of System 1, like those of System 2, are not restricted to the processing of current stimulation. Intuitive judgments deal with concepts as well as with percepts and

Figure 1
Process and Content in Two Cognitive Systems

	PERCEPTION	INTUITION SYSTEM 1	REASONING SYSTEM 2
PROCESS	Fast Parallel Automatic Effortless Associative Slow-learning Emotional		Slow Serial Controlled Effortful Rule-governed Flexible Neutral
CONTENT	Percepts Current stimulation Stimulus-bound	Conceptual representations Past, Present and Future Can be evoked by language	

can be evoked by language. In the model that is presented here, the perceptual system and the intuitive operations of System 1 generate *impressions* of the attributes of objects of perception and thought. These impressions are neither voluntary nor verbally explicit. In contrast, *judgments* are always intentional and explicit even when they are not overtly expressed. Thus, System 2 is involved in all judgments, whether they originate in impressions or in deliberate reasoning. The label *intuitive* is applied to judgments that directly reflect impressions—they are not modified by System 2.

As in several other dual-process models, one of the functions of System 2 is to monitor the quality of both mental operations and overt behavior (Gilbert, 2002; Stanovich & West, 2002). As expected for an effortful operation, the self-monitoring function is susceptible to dual-task interference. People who are occupied by a demanding mental activity (e.g., attempting to hold in mind several digits) are more likely to respond to another task by blurting out whatever comes to mind (Gilbert, 1989). The anthropomorphic phrase “System 2 monitors the activities of System 1” is used here as shorthand for a hypothesis about what would happen if the operations of System 2 were disrupted.

Kahneman and Frederick (2002) suggested that the monitoring is normally quite lax and allows many intuitive judgments to be expressed, including some that are erroneous. Shane Frederick (personal communication, April 29, 2003) has used simple puzzles to study cognitive self-monitoring, as in the following example: “A bat and a ball cost \$1.10 in total. The bat costs \$1 more than the ball. How much does the ball cost?” Almost everyone reports an initial tendency to answer “10 cents” because the sum \$1.10 separates naturally into \$1 and 10 cents and because 10 cents is about the right magnitude. Frederick found that many intelligent people yield to this immediate impulse: Fifty percent (47/93) of Princeton students and 56% (164/293) of students at the University of Michigan gave the wrong answer. Clearly, these respondents offered a response without checking it. The surprisingly high rate of errors in this easy problem illustrates how lightly the output of System 1 is monitored by System 2: People are not accustomed to thinking hard and are often content to trust a plausible judgment that quickly comes to mind. Remarkably, errors in this puzzle and in others of the same type were significant predictors of intolerance of delay and also of cheating behavior.

In the examples discussed so far, intuition was associated with poor performance, but intuitive thinking can also be powerful and accurate. High skill is acquired by prolonged practice, and the performance of skills is rapid and effortless. The proverbial master chess player who walks past a game and declares, “White mates in three,” without slowing is performing intuitively (Simon & Chase, 1973), as is the experienced nurse who detects subtle signs of impending heart failure (Gawande, 2002; Klein, 1998). Klein (2003, chapter 4) has argued that skilled decision makers often do better when they trust their intuitions than when they engage in detailed analysis. In the same vein,

Wilson and Schooler (1991) described an experiment in which participants who chose a poster for their own use were happier with it if their choice had been made intuitively than if it had been made analytically.

The Accessibility Dimension


A core property of many intuitive thoughts is that under appropriate circumstances, they come to mind spontaneously and effortlessly, like percepts. To understand intuition, then, one must understand why some thoughts come to mind more easily than others, why some ideas arise effortlessly and others demand work. The central concept of the present analysis of intuitive judgments and preferences is *accessibility*—the ease (or effort) with which particular mental contents come to mind. The accessibility of a thought is determined jointly by the characteristics of the cognitive mechanisms that produce it and by the characteristics of the stimuli and events that evoke it.

The question of why particular ideas come to mind at particular times has a long history in psychology. Indeed, this was the central question that the British empiricists sought to answer with laws of association. The behaviorists similarly viewed the explanation of “habit strength” or “response strength” as the main task of psychological theory, to be solved by a formulation integrating multiple determinants in the history and in the current circumstances of the organism. During the half century of the cognitive revolution, the measurement of reaction time became widely used as a general-purpose measure of response strength, and major advances were made in the study of why thoughts become accessible—notably, the distinctions between automatic and controlled processes and between implicit and explicit measures of memory. But no general concept was adopted, and research on the problem remained fragmented in multiple paradigms, variously focused on automaticity, Stroop interference, involuntary and voluntary attention, and priming.

Because the study of intuition requires a common concept, I adopt the term *accessibility*, which was proposed in the context of memory research (Tulving & Pearlstone, 1966) and of social cognition (Higgins, 1996) and is applied here more broadly than it was by these authors. In the present usage, the different aspects and elements of a situation, the different objects in a scene, and the different attributes of an object—all can be more or less accessible. Moreover, the determinants of accessibility subsume the notions of stimulus salience, selective attention, specific training, associative activation, and priming.

For an illustration of differential accessibility, consider Figures 2A and 2B. As one looks at the object in Figure 2A, one has immediate impressions of the height of the tower, the area of the top, and perhaps the volume of the tower. Translating these impressions into units of height or volume requires a deliberate operation, but the impressions themselves are highly accessible. For other attributes, no perceptual impression exists. For example, the total area that the blocks would cover if the tower were dismantled is not perceptually accessible, though it can be estimated by a deliberate procedure, such as multiplying the area of the

Figure 2
The Selective Accessibility of Natural Assessments



side of a block by the number of blocks. Of course, the situation is reversed with Figure 2B. Now, the blocks are laid out, and an impression of total area is immediately accessible, but the height of the tower that could be constructed with these blocks is not.

Some relational properties are accessible. Thus, it is obvious at a glance that Figures 2A and 2C are different but also that they are more similar to each other than either is to Figure 2B. Some statistical properties of ensembles are accessible, whereas others are not. For an example, consider the question "What is the average length of the lines in Figure 3?" This question is easily answered. When a set of objects of the same general kind is presented to an observer—whether simultaneously or successively—a representation of the set is computed automatically; this representation includes accurate impressions of the average (Ariely, 2001; Chong & Treisman, 2003). The representation of the prototype is highly accessible, and it has the character of a percept: One forms an impression of the typical line without choosing to do so. The only role for System 2 in this task is to map this impression of typical length onto the appropriate scale. In contrast, the answer to the question "What is the total length of the lines in the display?" does not come to mind without considerable effort.


These perceptual examples serve to establish the dimension of accessibility. At one end of this dimension are found operations that have the characteristics of perception and of the intuitive System 1: They are rapid, automatic, and effortless. At the other end are slow, serial, and effortful operations that people need a special reason to undertake. Accessibility is a continuum, not a dichotomy, and some effortful operations demand more effort than others. The acquisition of skill selectively increases the accessi-

bility of useful responses and of productive ways to organize information. The master chess player does not see the same board as the novice, and the skill of visualizing the tower that could be built from an array of blocks could surely be improved by prolonged practice.

Determinants of Accessibility

What becomes accessible in any particular situation is mainly determined, of course, by the actual properties of the object of judgment: It is easier to see a tower in Figure 2A than in Figure 2B because the tower in the latter is only virtual. Physical salience also determines accessibility: If a

Figure 3
The Selective Accessibility of Prototypical (Average) Features




large green letter and a small blue letter are shown at the same time, green will come to mind first. However, salience can be overcome by deliberate attention: An instruction to look for the smaller letter will enhance the accessibility of all its features, including its color. Motivationally relevant and emotionally arousing stimuli spontaneously attract attention. All the features of an arousing stimulus become accessible, including those that have no motivational or emotional significance. This fact is known, of course, to the designers of billboards.

The perceptual effects of salience and of spontaneous and voluntary attention have counterparts in the processing of more abstract stimuli. For example, the statements "Team A beat Team B" and "Team B lost to Team A" convey the same information. Because each sentence draws attention to its subject, however, the two versions make different thoughts accessible. Accessibility also reflects temporary states of priming and associative activation, as well as enduring operating characteristics of the perceptual and cognitive systems. For example, the mention of a familiar social category temporarily increases the accessibility of the traits associated with the category stereotype, as indicated by a lowered threshold for recognizing manifestations of these traits (Higgins, 1996; for a review, see Fiske, 1998). Moreover, the "hot" states of high emotional and motivational arousal greatly increase the accessibility of thoughts that relate to the immediate emotion and current needs, as well as reducing the accessibility of other thoughts (Loewenstein, 1996).

Some attributes, which Tversky and Kahneman (1983) called *natural assessments*, are routinely and automatically registered by the perceptual system or by System 1 without intention or effort. Kahneman and Frederick (2002) compiled a list of natural assessments with no claim to completeness. In addition to physical properties such as size, distance, and loudness, the list includes more abstract properties such as similarity (see, e.g., Tversky & Kahneman, 1983), causal propensity (Heider, 1944; Kahneman & Varey, 1990; Michotte, 1963), surprisingness (Kahneman & Miller, 1986), affective valence (see, e.g., Bargh, 1997; Cacioppo, Priester, & Berntson, 1993; Kahneman, Ritov, & Schkade, 1999; Slovic, Finucane, Peters, & MacGregor, 2002; Zajonc, 1980), and mood (Schwarz & Clore, 1983). Accessibility itself is a natural assessment—the routine evaluation of cognitive fluency in perception and memory (see, e.g., Jacoby & Dallas, 1981; Johnston, Dark, & Jacoby, 1985; Schwarz & Vaughn, 2002; Tversky & Kahneman, 1973).¹

The evaluation of stimuli as good or bad is a particularly important natural assessment. The evidence, both behavioral (Bargh, 1997; Zajonc, 1998) and neurophysiological (see, e.g., LeDoux, 2000), is consistent with the idea that the assessment of whether objects are good (and should be approached) or bad (and should be avoided) is carried out quickly and efficiently by specialized neural circuitry. Several authors have commented on the influence of this primordial evaluative system (here included in System 1) on the attitudes and preferences that people adopt con-

Figure 4
An Effect of Context on the Accessibility of Interpretations (After Bruner & Minturn, 1955)



sciously and deliberately (Epstein, 2003; Kahneman et al., 1999; Slovic et al., 2002; Wilson, 2000; Zajonc, 1998).

Figure 4 illustrates the effect of context on accessibility. An ambiguous stimulus that is perceived as a letter in a context of letters is seen as a number in a context of numbers. The figure also illustrates another point: The ambiguity is suppressed in perception. This aspect of the demonstration is spoiled for the reader who sees the two versions in close proximity, but when the two lines are shown separately, observers do not spontaneously become aware of the alternative interpretation. They "see" the interpretation that is the most likely in its context but have no subjective indication that it could be seen differently. Similarly, in bi-stable pictures such as the mother/daughter figure or the Necker cube, there is no perceptual representation of the instability. Almost no one (for a report of a tantalizing exception, see Wittreich, 1961) is able to see the Ames room as anything but rectangular, even when fully informed that the room is distorted and that the photograph does not provide enough information to specify its true shape. As the transactionalists who built the Ames room emphasized, perception is a choice of which people are not aware, and people perceive what has been chosen.

Uncertainty is poorly represented in intuition, as well as in perception. Indeed, the concept of judgment heuristics was invented to accommodate the observation that intuitive judgments of probability are mediated by attributes such as similarity and associative fluency, which are not intrinsically related to uncertainty. The central finding in studies of intuitive decisions, as described by Klein (1998), is that experienced decision makers working under pressure, such as captains of firefighting companies, rarely need to choose between options because in most cases only a single option comes to their mind. The options that were rejected are not

¹ The availability heuristic is based on an assessment of accessibility in which frequencies or probabilities are judged by the ease with which instances come to mind. Tversky and I were responsible for this terminological confusion (Tversky & Kahneman, 1973).

represented. Doubt is a phenomenon of System 2, a metacognitive appreciation of one's ability to think incompatible thoughts about the same thing.

Close counterfactual alternatives to what happened are perceived—one can see a horse that was catching up at the finish as almost winning the race (Kahneman & Varey, 1990). Norm theory (Kahneman & Miller, 1986) proposes that events evoke their own norms and that counterfactual alternatives to surprising occurrences are automatically accessible. In contrast to counterfactual alternatives to reality, competing interpretations of reality suppress each other: One does not see each horse in a close finish as both winning and losing.

As this discussion illustrates, much is known about the determinants of accessibility, but there is no general theoretical account of accessibility and no prospect of one emerging soon. In the context of research in judgment and decision making, however, the lack of a theory does little damage to the usefulness of the concept. In this respect, the conceptual status of the principles of accessibility resembles that of Gestalt principles of perceptual grouping, which are often invoked, both implicitly and explicitly, in the planning of research and in the interpretation of results. For these purposes, what matters is that empirical generalizations about the determinants of differential accessibility are widely accepted and that there are accepted procedures for testing the validity of particular hypotheses. For example, the claims about the differential accessibility of attributes in Figures 2 and 3 appeal to the consensual judgments of perceivers, but propositions about accessibility are also testable in other ways. In particular, judgments of relatively inaccessible properties are expected to be substantially slower and more susceptible to interference by concurrent mental activity, in comparison to judgments of accessible attributes.

Framing Effects

In Figure 2, the same property (the total height of a set of blocks) is highly accessible in one display and not so in another, although both displays contain the same information. This observation is entirely unremarkable—it does not seem shocking that some attributes of a stimulus are automatically perceived while others must be computed or that the same attribute is perceived in one display of an object but must be computed in another. In the context of decision making, however, similar observations raise a significant challenge to the rational-agent model. The assumption that preferences are not affected by variations of irrelevant features of options or outcomes has been called *extensionality* (Arrow, 1982) and *invariance* (Tversky & Kahneman, 1986); it is an essential aspect of the concept of rationality held in economic theory. Invariance is violated in demonstrations of *framing effects* such as the Asian disease problem (Tversky & Kahneman, 1981).

Problem 1—The Asian Disease

Imagine that the United States is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed.

Assume that the exact scientific estimates of the consequences of the programs are as follows:

If Program A is adopted, 200 people will be saved.

If Program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved.

Which one of the two programs would you favor?

In this version of the problem, a substantial majority of respondents favor Program A, indicating risk aversion. Other respondents, selected at random, receive a question in which the same cover story is followed by a different description of the options:

If Program A' is adopted, 400 people will die.

If Program B' is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die.

A clear majority of respondents now favor Program B', the risk-seeking option. Although there is no substantive difference between the versions, they evidently evoke different associations and evaluations. This is easiest to see in the certain option because outcomes that are certain are overweighted relative to outcomes of high or intermediate probability (Kahneman & Tversky, 1979). Thus, the certainty of saving people is disproportionately attractive, and the certainty of deaths is disproportionately aversive. These immediate affective responses respectively favor Program A over Program B and Program B' over Program A'. As in Figures 2A and 2B, the different representations of the outcomes highlight some features of the situation and mask others.

The question of how to determine whether two decision problems are the same or different does not have a general answer. To avoid this issue, Tversky and I restricted the definition of framing effects to discrepancies between choice problems that decision makers, upon reflection, consider effectively identical. The Asian disease problem passes this test: Respondents who are asked to compare the two versions almost always conclude that the same action should be taken in both. Observers agree that it would be frivolous to let a superficial detail of formulation determine a choice that has life-and-death consequences.

In another famous demonstration of an embarrassing framing effect, McNeil, Pauker, Sox, and Tversky (1982) induced different choices between surgery and radiation therapy by describing outcome statistics in terms of survival rates or mortality rates. Because 90% short-term survival is less threatening than 10% immediate mortality, the survival frame yielded a substantially higher preference for surgery. The framing effect was as pronounced among experienced physicians as it was among patients.

A different type of framing effect was demonstrated by Shafir (1993), who presented respondents with problems in which they played the role of a judge in adjudicating the custody of a child between divorcing parents. Each parent was described by a list of attributes. One of the descriptions

was richer than the other: It contained more negative and more positive attributes. The framing of the instruction was varied. Some respondents were asked which custody request should be accepted; others decided which request should be rejected. The rich description was selected under both instructions, presumably because its numerous advantages were salient (accessible) when the task was to choose which custody request to accept and its numerous disadvantages were salient when the focus of the task was rejection.

A large-scale study by LeBoeuf and Shafir (2003) examined an earlier claim that framing effects are reduced, in a between-participants design, for participants with high scores on "need for cognition" (Smith & Levin, 1996). The original effect was not replicated in the more extensive study. However, LeBoeuf and Shafir showed that more thoughtful individuals do show greater consistency in a within-participant design, where each respondent encounters both versions of each problem. This result is in accord with the present analysis. Respondents characterized by an active System 2 are more likely than others to notice the relationship between the two versions and to ensure the consistency of the responses to them. Thoughtfulness confers no advantage in the absence of a usable cue and is therefore irrelevant to performance in the between-participants design. As was noted earlier, the accessibility of a thought depends both on the characteristics of the cognitive system and on the presence of an appropriate stimulus.

Framing effects are not restricted to decision making: Simon and Hayes (1976) documented an analogous observation in the domain of problem solving. They constructed a collection of transformation puzzles, all formally identical to the tower of Hanoi problem, and found that these "problem isomorphs" varied greatly in difficulty. For example, the initial state and the target state were described in two of the versions as three monsters holding balls of different colors. The state transitions were described in one version as changes in the color of the balls and in the other as balls being passed from one monster to another. The puzzle was solved much more easily when framed in terms of motion. The authors commented that "it would be possible for a subject to seek that representation which is simplest, according to some criterion, or to translate all such problems into the same, canonical, representation" but that "subjects will not employ such alternative strategies, even though they are available, but will adopt the representation that constitutes the most straightforward translation" (Simon & Hayes, 1976, p. 183).

The basic principle of framing is the passive acceptance of the formulation given. This general principle applies equally as well to puzzles, to the displays of Figure 2, and to the standard framing effects. People do not spontaneously compute the height of a tower that could be built from an array of blocks, and they do not spontaneously transform the representation of puzzles or decision problems. The brain mechanisms that support the comprehension of language have a substantial ability to strip the surface details and get to the gist of meaning in an utterance, but this ability is limited as well. Few people are able

to recognize 137×24 and 3,288 as the same number without going through some elaborate computations. Invariance cannot be achieved by a finite mind.

The impossibility of invariance raises significant doubts about the descriptive realism of rational-choice models (Tversky & Kahneman, 1986). Absent a system that reliably generates appropriate canonical representations, intuitive decisions are shaped by the factors that determine the accessibility of different features of the situation. Highly accessible features influence decisions, whereas features of low accessibility are largely ignored. Unfortunately, there is no reason to believe that the most accessible features are also the most relevant to a good decision.

Changes or States: Prospect Theory


A general property of perceptual systems is that they appear designed to enhance the accessibility of changes and differences (Palmer, 1999). Perception is *reference dependent*: The perceived attributes of a focal stimulus reflect the contrast between that stimulus and a context of prior and concurrent stimuli. Figure 5 illustrates reference dependence in vision. The two enclosed squares have the same luminance, but they do not appear equally bright. The point of the demonstration is that the brightness of an area is not a single-parameter function of the light energy that reaches the eye from that area. An account of perceived brightness also requires a parameter for a reference value (often called *adaptation level*), which is influenced by the luminance of neighboring areas.

The reference value to which current stimulation is compared also reflects the history of adaptation to prior stimulation. A familiar demonstration involves three buckets of water of different temperatures, arranged from cold on the left to hot on the right, with tepid in the middle. In the adapting phase, the left and right hands are immersed in cold and hot water, respectively. The initially intense sensations of cold and heat gradually wane. When both hands are then immersed in the middle bucket, the experience is heat in the left hand and cold in the right hand.

Reference Dependence in Choice

The facts of perceptual adaptation were in our minds when Tversky and I began our joint research on decision making under risk. Guided by the analogy of perception, we expected the evaluation of decision outcomes to be reference dependent. We noted, however, that reference dependence is incompatible with the standard interpretation of expected utility theory, the prevailing theoretical model of risky choice. This deficiency can be traced to the brilliant essay that introduced the first version of that theory (Bernoulli, 1738/1954). Bernoulli's great innovation was to abandon the standard way of evaluating gambles by their expected value—the weighted average of their outcomes (in ducats), each weighted by its probability. Instead, Bernoulli proposed that the value of a gamble is the probability-weighted average of the psychological values (utilities) of its outcomes, which he defined as states of wealth. Developing an

Figure 5
Simultaneous Contrast and Reference Dependence



argument that anticipated the psychophysics of Weber and Fechner by more than a century, Bernoulli concluded that the utility function of wealth is logarithmic. Economists discarded the logarithmic function long ago, but the idea that decision makers evaluate outcomes by the utility of wealth positions has been retained in economic analyses for almost 300 years. This is rather remarkable because the idea is easily shown to be wrong; I call it Bernoulli's error.

Bernoulli's (1738/1954) model of utility is flawed because it is *reference independent*: It assumes that the utility that is assigned to a given state of wealth does not vary with the decision maker's initial state of wealth. This assumption flies against a basic principle of perception, where the effective stimulus is not the new level of stimulation but the difference between it and the existing adaptation level. The analogy to perception suggests that the carriers of utility are likely to be gains and losses rather than states of wealth, and this suggestion is amply supported by the evidence of both experimental and observational studies of choice (see Kahneman & Tversky, 2000). The present discussion relies on two thought experiments of the kind that Tversky and I devised in the process of developing the model of risky choice that we called *prospect theory* (Kahneman & Tversky, 1979).

Problem 2

Would you accept this gamble?

50% chance to win \$150

50% chance to lose \$100

Would your choice change if your overall wealth were lower by \$100?

There will be few takers of the gamble in Problem 2. The experimental evidence shows that most people reject a gamble with even chances to win and lose unless the possible win is at least twice the size of the possible loss

(see, e.g., Tversky & Kahneman, 1992). The answer to the second question is, of course, negative.

Next, consider Problem 3.

Problem 3

Which would you choose?

Lose \$100 with certainty

or

50% chance to win \$50

50% chance to lose \$200

Would your choice change if your overall wealth were higher by \$100?

In Problem 3, the gamble appears much more attractive than the sure loss. Experimental results indicate that risk-seeking preferences are held by a large majority of respondents in choices of this kind (Kahneman & Tversky, 1979). Here again, the idea that a change of \$100 in total wealth would affect preferences cannot be taken seriously.

Problems 2 and 3 evoke sharply different preferences, but from a Bernoullian perspective, the difference is a framing effect: When stated in terms of final wealth, the problems only differ in that all values are lower by \$100 in Problem 3—surely, an inconsequential variation. Tversky and I examined many choice pairs of this type early in our explorations of risky choice and concluded that the abrupt transition from risk aversion to risk seeking could not plausibly be explained by a utility function for wealth. Preferences appeared to be determined by attitudes to gains and losses, defined relative to a reference point, but Bernoulli's (1738/1954) theory and its successors did not incorporate a reference point. We therefore proposed an alternative theory of risk in which the carriers of utility are gains and losses—changes of wealth rather than states of wealth. Prospect theory (Kahneman & Tversky, 1979) em-

braces the idea that preferences are reference dependent and includes the extra parameter that is required by this assumption.

The distinctive predictions of prospect theory follow from the shape of the value function, which is shown in Figure 6. The value function is defined on gains and losses and is characterized by four features: (a) It is concave in the domain of gains, favoring risk aversion; (b) it is convex in the domain of losses, favoring risk seeking; (c) most important, the function is sharply kinked at the reference point and *loss averse*—steeper for losses than for gains by a factor of about 2–2.5 (Kahneman, Knetsch, & Thaler, 1991; Tversky & Kahneman, 1992); and (d) several studies suggest that the functions in the two domains are fairly well approximated by power functions with similar exponents, both less than unity (Swalm, 1966; Tversky & Kahneman, 1992). The power function is not surprising because the value function is a psychophysical mapping. However, the value function is not expected to describe preferences for losses that are large relative to total assets, where ruin or near ruin is a possible outcome.

Bernoulli's error—the assumption that the carriers of utility are final states—is not restricted to decision making under risk. Indeed, the error of reference independence is built into the standard representation of indifference maps, a basic tool of economic thinking. It is puzzling to a psychologist that these maps do not include a representation of the decision maker's current holdings of various goods—the counterpart of the reference point in prospect theory. The parameter is not included, of course, because economic theory assumes that it does not matter.

The core idea of prospect theory—that the value function is kinked at the reference point and loss averse—became useful to economics when Thaler (1980) used it to explain riskless choices. In particular, loss aversion explained a violation of consumer theory that Thaler identified and labeled the *endowment effect*: The maximum amount that people pay to acquire a good is commonly much less than the minimal amount they demand to part

from it once they own it. The selling price often exceeds the buying price by a factor of 2 or more (Kahneman, Knetsch, & Thaler, 1990, 1991; Tversky & Kahneman, 1991). The interpretation is straightforward: A good is worth more when it is considered as something that could be lost or given up than when it is evaluated as a potential gain.

Reference dependence and loss aversion help account for several phenomena of choice. The familiar observation that out-of-pocket losses are much more distressing than foregone gains is readily explained if these outcomes are evaluated on different limbs of the value function. The distinction between actual losses and lost opportunities is recognized in applications of the law (Cohen & Knetsch, 1992) and in lay intuitions about rules of fairness in the market (Kahneman, Knetsch, & Thaler, 1986). Loss aversion also contributes to the well-documented status-quo bias (Samuelson & Zeckhauser, 1988). Because the reference point is usually the status quo, the properties of alternative options are evaluated as advantages or disadvantages relative to the current situation, and the disadvantages of the alternatives loom larger than their advantages. Other applications of the concept of loss aversion are documented in several chapters in Kahneman and Tversky (2000).

Narrow Framing

The idea that the carriers of utility are changes of wealth rather than asset positions was described as the cornerstone of prospect theory (Kahneman & Tversky, 1979, p. 273). This statement implies that choices are always made by considering gains and losses rather than final states, but there are exceptions to this claim. For an example, consider Problem 4.

Problem 4

Please estimate your total wealth, call it W .

Which of these situations is more attractive:

You own W

or


50% chance that you own $W - \$100$

50% chance that you own $W + \$150$

Informal experiments with problems of this type have consistently yielded a mild preference for the uncertain state of wealth and a strong impression that the stakes mentioned in the question are entirely negligible.

In terms of final states of wealth, Problem 4 is identical to Problem 2. Furthermore, most respondents agree, upon reflection, that the difference between the problems is inconsequential—too slight to justify different choices. Thus, the discrepant preferences observed in these two problems satisfy the definition of a framing effect. The manipulation of accessibility that produces this framing effect is straightforward. The gamble of Problem 2 is likely to evoke an evaluation of the emotions associated with the immediate outcomes, and the formulation does not bring to

Figure 6
The Value Function of Prospect Theory



mind thoughts of overall wealth. In contrast, the formulation of Problem 4 favors a view of the uncertainty as trivially small in relation to *W* and does not evoke the emotional asymmetry of gains or losses. The two problems elicit different representations and, therefore, different preferences. Indeed, they are explained by different theories. Prospect theory (where value is attached to changes) is not applicable to Problem 4, and standard utility theory (where utility is attached to wealth) is not applicable to Problem 2.

Some real-world choices are made in the wealth frame. In particular, financial advisors and decision analysts often insist on formulating outcomes in terms of assets when eliciting their clients' preferences. These cases are rare, however. The effective outcomes in the overwhelming majority of decision problems are gains and losses, and Bernoulli's formulation is not useful to explain risky choices that are so framed. There is a genuine puzzle here: Why has a transparently incorrect model been retained for so long as the dominant theory of choice in economics and in other social sciences? The answer may well be that the assignment of utility to final states is compatible with the general assumption of rationality in economic theorizing.

Consider Problem 5.

Problem 5

Two persons get their monthly report from a broker:

A is told that her wealth went from 4M to 3M.

B is told that her wealth went from 1M to 1.1M.

- (i) Who of the two individuals has more reason to be satisfied with her financial situation?
- (ii) Who is happier today?

Problem 5 highlights the contrasting interpretations of utility in theories that define outcomes as states or as changes. In Bernoulli's analysis, only the first of the two questions is relevant, and only long-term consequences matter. Prospect theory, in contrast, is concerned with short-term outcomes, and the value function presumably reflects an anticipation of the valence and intensity of the emotions that are experienced at moments of transition from one state to another (Kahneman, 2000b, 2000c; Mellers, 2000). Which of these concepts of utility is more useful? For descriptive purposes, the more myopic notion is superior, but the prescriptive norms of reasonable decision making favor the long-term view. The Bernoullian definition of relevant outcomes is a good fit in a rational-agent model.

A particularly unrealistic implication of the rational-agent model is that agents make their choices in a comprehensively inclusive context, which incorporates all the relevant details of the present situation as well as expectations about all future opportunities and risks. Much evidence supports the contrasting claim that people's views of decisions and outcomes are normally characterized by *narrow framing* (Kahneman & Lovallo, 1993) and by the related notion of *mental accounting* (Thaler, 1985, 1999). The

prevalence of the gains/losses frame illustrates narrow framing.

For another example, note that it has appeared natural to consider each of the preceding choice problems on its own, as a separate decision. However, this framing is unreasonable for decision makers who expect to live long enough to make many other decisions about accepting gambles. In the broader view, the choice of the moment is an occasion to apply a general policy, not an occasion for an isolated decision. Several experiments have shown that people are much more willing to accept gambles of positive expected value if they are assured of the opportunity to play several times (Benartzi & Thaler, 1999; Kahneman & Lovallo, 1993; Keren & Wagenaar 1987; Tversky & Redelmeier, 1992). In the broad view, of course, the distinction between single play and multiple play is largely irrelevant because life is likely to provide additional opportunities to gamble. Here again, however, people accept the frames that are suggested to them: They consider repeated plays when instructed to do so but focus on a single problem when it is presented on its own. A shared feature of these examples is that decisions made in narrow frames depart far more from risk neutrality than decisions made in a more inclusive context.

The prevalence of narrow frames is an effect of accessibility, which can be understood by referring to the displays of blocks in Figure 2. The same set of blocks is framed as a tower in Figure 2A and as a flat array in Figure 2B. Although it is possible to see a tower in Figure 2B, it is much easier to do so in Figure 2A. Narrow frames generally reflect the structure of the environment in which decisions are made. The choices that people face arise one at a time, and the principle of passive acceptance suggests that they are considered as they arise. The problem at hand and the immediate consequences of the choice are far more accessible than all other considerations, and as a result, decision problems are framed far more narrowly than the rational model assumes.

It is worth noting that an exclusive concern with the broad view and with the long term may be prescriptively sterile because the long term is not where life is lived. Utility cannot be divorced from emotion, and emotion is triggered by changes. A theory of choice that completely ignores feelings such as the pain of losses and the regret of mistakes is not just descriptively unrealistic. It also leads to prescriptions that do not maximize the utility of outcomes as they are actually experienced—that is, utility as Bentham conceived it (Kahneman, 1994, 2000a; Kahneman, Wakker, & Sarin, 1997).

Attribute Substitution: A Model of Judgment by Heuristic

The first joint research program that Tversky and I undertook was a study of various types of judgment about uncertain events, including numerical predictions and assessments of the probabilities of hypotheses. We reviewed this work in an integrative article (Tversky & Kahneman, 1974), which aimed to show

that people rely on a limited number of heuristic principles which reduce the complex tasks of assessing probabilities and predicting values to simpler judgmental operations. In general, these heuristics are quite useful, but sometimes they lead to severe and systematic errors. (Tversky & Kahneman, 1974, p. 1124)

The second paragraph of that article introduced the idea that "the subjective assessment of probability resembles the subjective assessments of physical quantities such as distance or size. These judgments are all based on data of limited validity, which are processed according to heuristic rules" (Tversky & Kahneman, 1974, p. 1124). The concept of *heuristic* was illustrated by the role of the blur of contours as a potent determinant of the perceived distance of mountains. The observation that reliance on blur as a distance cue causes distances to be overestimated on foggy days and underestimated on clear days was the example of a heuristic-induced *bias*. As this example illustrates, heuristics of judgment were to be identified by the characteristic errors that they tend to cause.

Three heuristics of judgment, labeled *representativeness*, *availability*, and *anchoring*, were described in the 1974 review, along with a dozen systematic biases, including nonregressive prediction, neglect of base-rate information, overconfidence, and overestimates of the frequency of events that are easy to recall. Some of the biases were identified by systematic errors in estimates of known quantities and statistical facts. Other biases were identified by systematic discrepancies between the regularities of intuitive judgments and the principles of probability theory, Bayesian inference, or regression analysis. The article launched the so-called heuristics and biases approach to the study of intuitive judgment, which has been the topic of a substantial research literature (Gilovich, Griffin, & Kahneman, 2002; Kahneman, Slovic, & Tversky, 1982) and has also been the focus of substantial controversy.

Shane Frederick and I recently revisited the conception of heuristics and biases in the light of developments in the study of judgment and in the broader field of cognitive psychology in the intervening three decades (Kahneman & Frederick, 2002). The new model departs from the original formulation of heuristics in three significant ways: (a) It proposes a common process of attribute substitution to explain how judgment heuristics work, (b) it extends the concept of heuristic beyond the domain of judgments about uncertain events, and (c) it includes an explicit treatment of the conditions under which intuitive judgments are modified or overridden by the monitoring operations associated with System 2.

Attribute Substitution

The 1974 article did not include a definition of judgmental heuristics. Heuristics were described at various times as principles, as processes, or as sources of cues for judgment. The vagueness did no damage because the research program focused on a total of three heuristics of judgment under uncertainty that were separately defined in adequate detail. In contrast, Kahneman and Frederick (2002) offered an explicit definition of a generic heuristic process of

attribute substitution: A judgment is said to be mediated by a heuristic when the individual assesses a specified *target attribute* of a judgment object by substituting a related *heuristic attribute* that comes more readily to mind. This definition elaborates a theme of the early research, namely, that people who are confronted with a difficult question sometimes answer an easier one instead. The word *heuristic* is used in two senses in the new definition. The noun refers to the cognitive process, and the adjective in *heuristic attribute* specifies the attribute that is substituted in a particular judgment. For example, the representativeness heuristic is the use of representativeness as a heuristic attribute to judge probability. The definition of heuristics by attribute substitution does not coincide perfectly with the original conception offered by Tversky and Kahneman (1974). In particular, the new concept excludes anchoring effects, in which judgment is influenced by temporarily raising the accessibility of a particular value of the target attribute, relative to other values of the same attribute.

For a perceptual example of attribute substitution, consider the question "What are the sizes of the two horses in Figure 7, as they are shown on the page?" The images are in fact identical in size, but the figure produces a compelling illusion. The target attribute that the observer is instructed to report is two-dimensional size, but the responses actually map an impression of three-dimensional size onto units of length that are appropriate to the required judgment. In the terms of the model, three-dimensional size is the heuristic attribute. As in other cases of attribute substitution, the illusion is caused by differential accessibility. An impression of three-dimensional size is the only impression of size that comes to mind for naïve observers—painters and experienced photographers are able to do better—and it produces a perceptual illusion in the judgment of picture size. The cognitive illusions that are produced by attribute substitution have the same character: An impression of one attribute is mapped onto the scale of another, and the judge is normally unaware of the substitution.

The most direct evidence for attribute substitution was reported by Kahneman and Tversky (1973) in a task of categorical prediction. There were three experimental groups in the experiment. Participants in a base-rate group evaluated the relative frequencies of graduate students in nine categories of specialization.² Mean estimates ranged from 20% for humanities and education to 3% for library science.

Two other groups of participants were shown the same list of areas of graduate specialization and the following description of a fictitious graduate student.

Tom W. is of high intelligence, although lacking in true creativity. He has a need for order and clarity, and for neat and tidy systems in which every detail finds its appropriate place. His writing is rather dull and mechanical, occasionally enlivened by somewhat

² The categories were business administration, computer science, engineering, humanities and education, law, library science, medicine, physical and life sciences, and social sciences and social work.

Figure 7

Attribute Substitution in Perception: A Highly Accessible Heuristic Attribute (Three-Dimensional Size) Substitutes for a Less Accessible Target Attribute (Picture Size)



Note. Photo by Lenore Shoham, 2003.

corny puns and by flashes of imagination of the sci-fi type. He has a strong drive for competence. He seems to have little feel and little sympathy for other people and does not enjoy interacting with others. Self-centered, he nonetheless has a deep moral sense.

Participants in a similarity group ranked the nine fields by the degree to which Tom W. “resembles a typical graduate student” (in that field). The description of Tom W. was deliberately constructed to make him more representative of the less populated fields, and this manipulation was successful: The correlation between the average representativeness rankings and the estimated base rates of fields of specialization was -0.62 . Participants in the probability group ranked the nine fields according to the likelihood that Tom W. would have specialized in each. The respondents in the latter group were graduate students in psychology at major universities. They were told that the personality sketch had been written by a psychologist when Tom W. was in high school, on the basis of personality tests of dubious validity. This information was intended to discredit the description as a source of valid information.


The statistical logic is straightforward. A description based on unreliable information must be given little weight, and predictions made in the absence of valid evidence must revert to base rates. This reasoning implies that judgments of probability should be highly correlated with the corresponding base rates in this problem.

The psychology of the task is also straightforward. The similarity of Tom W. to various stereotypes is a highly accessible natural assessment, whereas judgments of probability are difficult. The respondents are therefore expected to substitute a judgment of similarity (representativeness) for the required judgment of probability. The two instructions—to rate similarity or probability—should therefore elicit similar judgments.

The scatter plot of the mean judgments of the two groups is presented in Figure 8A. As the figure shows, the correlation between judgments of probability and similarity is nearly perfect (0.98). The correlation between judgments of probability and base-rates is -0.63 . The results are in perfect accord with the hypothesis of attribute substitution.

Figure 8

Mean Judgments of Probability Are Plotted Against Mean Judgments of Similarity (Representativeness) for Eight Possible Outcomes in the Linda Problem



They also confirm a bias of *base-rate neglect* in this prediction task.

Figure 8B shows the results of another study in the same design, in which respondents were shown the description of a woman named Linda and a list of eight possible outcomes describing her present employment and activities. The two critical items in the list were number 6 ("Linda is a bank teller") and the conjunction item, number 8 ("Linda is a bank teller and active in the feminist movement"). The other six possibilities were unrelated and miscellaneous (e.g., elementary school teacher, psychiatric social worker). As in the Tom W. problem, some respondents were required to rank the eight outcomes by the similarity of Linda to the category prototypes; others ranked the same outcomes by probability.

Linda is 31 years old, single, outspoken and very bright. She majored in philosophy. As a student she was deeply concerned with issues of discrimination and social justice and also participated in antinuclear demonstrations.

As might be expected, 85% of respondents in the similarity group ranked the conjunction item (number 8) higher than its constituent, indicating that Linda resembles the image of a feminist bank teller more than she resembles a stereotypical bank teller. This ordering of the two items is quite reasonable for judgments of similarity. However, it is much more problematic that 89% of respondents in the probability group also ranked the conjunction higher than its constituent. This pattern of probability judgments violates monotonicity and has been called the *conjunction fallacy* (Tversky & Kahneman, 1983).

The results shown in Figure 8 are especially compelling because the responses were rankings. The large variability of the average rankings of both attributes indicates highly consensual responses and nearly total overlap in the systematic variance. Stronger support for attribute substi-

tution could hardly be imagined. Other tests of representativeness in the heuristic elicitation design have been equally successful (Bar-Hillel & Neter, 2002; Tversky & Kahneman, 1982). The same design was also applied extensively in studies of support theory (Tversky & Koehler, 1994; for a review, see Brenner, Koehler, & Rottenstreich, 2002). In one of the studies reported by Tversky and Koehler (1994), participants rated the probability that the home team would win in each of 20 specified basketball games and later provided ratings of the relative strength of the two teams, using a scale in which the strongest team in the tournament was assigned a score of 100. The correlation between normalized strength ratings and judged probabilities was 0.99.

The essence of attribute substitution is that respondents offer a reasonable answer to a question that they have not been asked. An alternative interpretation that must be considered is that the respondents' judgments reflect their understanding of the question that was posed. This may be true in some situations: It is not unreasonable to interpret a question about the probable outcome of a basketball game as referring to the relative strength of the competing teams. In many other situations, however, attribute substitution occurs even when the target and heuristic attributes are clearly distinct. For example, it is highly unlikely that educated respondents have a concept of probability that coincides precisely with similarity or that they are unable to distinguish picture size from object size. A more plausible hypothesis is that an evaluation of the heuristic attribute comes immediately to mind and that its associative relationship with the target attribute is sufficiently close to pass the permissive monitoring of System 2. Respondents who substitute one attribute for another are not confused about the question that they are trying to answer—they simply fail to notice that they are answering a different one. When they do notice the discrepancy and suspect a bias, they

either modify the intuitive judgment or abandon it altogether.

As illustrated by its use in the interpretation of the visual illusion of Figure 7, the definition of judgment heuristics by the mechanism of attribute substitution applies to many situations in which people make a judgment that is not the one they intended to make. There is no finite list of heuristic attributes. Kahneman and Frederick (2002) illustrated this conception by a study by Strack, Martin, and Schwarz (1988) in which college students answered a survey that included these two questions: "How happy are you with your life in general?" and "How many dates did you have last month?" The correlation between the two questions was negligible when they occurred in the order shown, but it rose to 0.66 when the dating question was asked first. The model of attribute substitution suggests that the dating question automatically evokes an affectively charged evaluation of one's satisfaction in that domain of life, which lingers to become the heuristic attribute when the happiness question is subsequently encountered. The underlying correlation between the judgment and the heuristic attribute is surely higher than the observed value of 0.66, which is attenuated by measurement error. The same experimental manipulation of question order was used in another study to induce the use of marital satisfaction as a heuristic attribute for well-being (Schwarz, Strack, & Mai, 1991). The success of these experiments suggests that ad hoc attribute substitution is a frequent occurrence. It is important to note that the present treatment does not make specific predictions about the heuristics that will be used in particular circumstances. It only provides (a) an approach that helps generate such predictions, based on the considerations of relative accessibility that were discussed earlier, and (b) two separate methods for testing heuristics, by examining predicted biases of judgment and by direct comparisons of the target and heuristic attributes.

The Affect Heuristic

The idea of an *affect heuristic* (Slovic et al., 2002) is probably the most important development in the study of judgment heuristics in the past few decades. There is compelling evidence for the proposition that every stimulus evokes an affective evaluation, which is not always conscious (see reviews by Bargh, 1997; Zajonc, 1980, 1998). Affective valence is a natural assessment and, therefore, a candidate for substitution in the numerous responses that express attitudes. Slovic and his colleagues (Slovic et al., 2002) discussed how a basic affective reaction can be used as the heuristic attribute for a wide variety of more complex evaluations, such as the cost-benefit ratio of technologies, the safe concentration of chemicals, and even the predicted economic performance of industries. Their treatment of the affect heuristic fits the present model of attribute substitution.

In the same vein, Kahneman and Ritov (1994) and Kahneman et al. (1999) proposed that an automatic affective valuation—the emotional core of an attitude—is the main determinant of many judgments and behaviors. In the study by Kahneman and Ritov, 37 public causes were

ranked by average responses to questions about (a) the importance of the issues, (b) the size of the donation that respondents were willing to make, (c) political support for interventions, and (d) the moral satisfaction associated with a contribution. The rankings were all very similar. In the terms of the present analysis, the same heuristic attribute (affective valuation) was mapped onto the distinct scales of different target attributes. Similarly, Kahneman, Schkade, and Sunstein (1998) interpreted jurors' assessments of punitive awards as a mapping of outrage onto a dollar scale of punishments. In an article titled "Risk as Feelings," Loewenstein, Weber, Hsee, and Welch (2001) offered a closely related analysis in which emotional responses, such as the intensity of fear, govern diverse judgments (e.g., ratings of the probability of a disaster).

In terms of the scope of responses that it governs, the natural assessment of affect should join representativeness and availability in the list of general-purpose heuristic attributes. The failure to identify the affect heuristic much earlier and its enthusiastic acceptance in recent years reflect significant changes in the general climate of psychological opinion. It is worth noting that the idea of purely cognitive biases appeared novel and distinctive in the early 1970s because the prevalence of motivated and emotional biases of judgment was taken for granted by the social psychologists of the time. There followed a period of intense emphasis on cognitive processes in psychology generally and in the field of judgment in particular. It took another 30 years to achieve what now appears to be a more integrated view of the role of affect in intuitive judgment.

The Accessibility of Corrective Thoughts

The present treatment assumes that System 2 continuously monitors the tentative judgments and intentions that System 1 produces. This assumption implies that errors of intuitive judgment involve failures of both systems: System 1, which generates the error, and System 2, which fails to detect and correct it (Kahneman & Tversky, 1982). To illustrate this point, Kahneman and Frederick (2002) revisited the visual example that Tversky and Kahneman (1974) had used to explain how heuristics generate biases: Blur is a good cue to the distance of mountains, but reliance on this cue causes predictable errors in estimates of distance on sunny or hazy days. The analogy was apt, but the analysis of the perceptual example neglected an important fact. Observers know, of course, whether the day is sunny or hazy. They could therefore apply this knowledge to counteract the bias—but unless they have been trained as sharpshooters, they are unlikely to do so. Contrary to what the early treatment implied, the use of blur as a cue does not inevitably lead to bias in the judgment of distance—the error could just as well be described as a failure to assign adequate negative weight to ambient haze. The effect of haziness on impressions of distance is a failing of System 1: The perceptual system is not designed to correct for this variable. The effect of haziness on judgments of distance is

a separate failure of System 2. The analysis extends readily to errors of intuitive judgment.

The observation that it is possible to design experiments in which cognitive illusions disappear has sometimes been used as an argument against the usefulness of the notions of heuristics and biases (see, e.g., Gigerenzer, 1991). In the present framework, however, there is no mystery about the conditions under which illusions appear or disappear: An intuitive judgment will be modified or overridden if System 2 identifies it as biased. This argument is not circular because a great deal is known about the conditions under which corrections will or will not be made and because hypotheses about the role of System 2 can be tested.

In the context of an analysis of accessibility, the question of when intuitive judgments will be corrected is naturally rephrased: When will corrective thoughts be sufficiently accessible to intervene in the judgment? There have been three lines of research on this issue. One explored the conditions that influence the general efficacy of System 2 and thereby the likelihood that potential errors will be detected and prevented. Other lines of research investigated the factors that determine the accessibility of relevant metacognitive knowledge and the accessibility of relevant statistical rules.

The corrective operations of System 2 are impaired by time pressure (Finucane, Alhakami, Slovic, & Johnson, 2000), by concurrent involvement in a different cognitive task (Gilbert, 1989, 1991, 2002), by performing the task in the evening for morning people and in the morning for evening people (Bodenhausen, 1990), and, surprisingly, by being in a good mood (Bless et al., 1996; Isen, Nygren, & Ashby, 1988). Conversely, the facility of System 2 is positively correlated with intelligence (Stanovich & West, 2002), with need for cognition (Shafir & LeBoeuf, 2002), and with exposure to statistical thinking (Agnoli, 1991; Agnoli & Krantz, 1989; Nisbett, Krantz, Jepson, & Kunda, 1983).

When people become aware of using a heuristic, they correct their judgment accordingly and may even overcorrect. For example, Schwarz and Clore (1983) showed that the normal effect of rainy weather on reports of general happiness is eliminated when respondents are first asked about the weather. The question about the weather has a metacognitive effect: It reminds respondents that they should not allow their judgments of well-being to be influenced by a transient weather-related mood. Schwarz, Bless, et al. (1991) and Oppenheimer (in press) showed similar discounting effects in studies of the availability heuristic. In an elegant series of experiments, Oppenheimer showed that respondents who were asked to estimate the frequency of surnames in the U.S. population even tended to underestimate the frequency of famous surnames, such as Bush, as well as the population frequency of their own surname. He also reported that an availability effect (overestimating the frequency of words that contain letters of one's initials) was replaced by a significant effect in the opposite direction when people were first required to write down their initials. It may be significant that these demonstrations of

metacognitive corrections were concerned with the availability heuristic and not with representativeness. The distinction between objective frequency and the availability of instances to memory is far more transparent than the distinction between probability and similarity, and it may be correspondingly easier to recognize availability biases in frequency judgments than to identify representativeness biases in statistical reasoning.

Nisbett, Krantz, and their colleagues mounted a substantial research program to investigate the factors that control the accessibility of *statistical heuristics*—rules of thumb that people can be trained to apply to relevant problems, such as “consider the size of the sample” (Nisbett et al., 1983). In one of their studies, Nisbett et al. (1983) compared formally identical problems that differed in content. They found that statistical reasoning was most likely to be evoked in the context of games of chance, was occasionally evoked in situations involving sports, but was relatively rare when the problems concerned the psychology of individuals. They also showed that the explicit mention of a sampling procedure facilitated statistical thinking (Nisbett et al., 1983; see also Gigerenzer, Hell, & Blank, 1988). Zukier and Pepitone (1984) found that respondents were more likely to use base-rate information when instructed to think as statisticians than when instructed to emulate psychologists. Agnoli and Krantz (1989) reported that brief training in the logic of sets improved performance in a simple version of the Linda problem. The findings indicate that the accessibility of statistical heuristics can be enhanced in at least three ways: by increasing the vigilance of the monitoring activities, by providing stronger cues to the relevant rules, and by extensive training in applied statistical reasoning.³

In the absence of primes and reminders, the accessibility of statistical heuristics is low. For an example, it is useful to consider how System 2 might have intervened in the problems of Tom W. and Linda that were described in an earlier section.

Tom W does look like a library science person, but there are many more graduate students in humanities and social sciences. I should adjust my rankings accordingly.

Linda cannot be more likely to be a feminist bank teller than to be a bank teller. I must rank these two outcomes accordingly.

Neither of these examples of reasoning exceeds the intellectual reach of the graduate students at major universities whose rankings were shown in Figure 8. However, the data indicate that very few respondents actually came up with these elementary corrections.

The present analysis of judgment implies that statistical training does not eradicate intuitive heuristics such as availability or representativeness but only enables people to avoid some biases under favorable circumstances. The results of Figure 8 support this prediction. In the absence of strong cues to remind them of their statistical knowledge,

³ The intervention of System 2 does not guarantee a correct response. The rules that people apply in deliberate reasoning are sometimes false.

statistically knowledgeable graduate students made categorical predictions like everybody else—by representativeness. However, statistical sophistication made a difference in a more direct version of the Linda problem, which required respondents to compare the probabilities of Linda being “a bank teller” or “a bank teller who is active in the feminist movement” (Tversky & Kahneman, 1983). The incidence of errors remained high for the statistically naïve even in that transparent version, but the error rate dropped dramatically among the sophisticated.

Analogous corrections can be expected for other intuitive judgments: For example, reports of subjective well-being are strongly influenced by current mood and current preoccupations (Schwarz & Strack, 1999), but reminding respondents to think broadly about their lives would certainly cause them to bring other relevant considerations to bear on their responses. Similarly, the initial punitive decisions of jurors are likely to reflect an outrage heuristic (Kahneman, Schkade, & Sunstein, 1998), but jurors can also be instructed to consider other factors.

The analysis of corrective thoughts has a significant methodological implication: Different research designs are appropriate for the study of System 1 and of System 2. If the goal of the research is to study intuitive judgment, the design should minimize the role of deliberation and self-critical reflection. Intuitive judgments and preferences are therefore best studied in between-participants designs and in short experiments that provide little information about the experimenter’s aims. Within-participant designs with multiple trials should be avoided because they encourage the participants to search for consistent strategies to deal with the task. Within-participant factorial designs are particularly undesirable because they provide an unmistakable cue that any factor that is varied systematically must be relevant to the target attribute (Kahneman & Frederick, 2002). The difficulties of these experimental designs were noted long ago by Kahneman and Tversky, who pointed out that “within-subject designs are associated with significant problems of interpretation in several areas of psychological research (Poulton, 1975). In studies of intuition, they are liable to induce the effect that they are intended to test” (Kahneman and Tversky, 1982, p. 500). Unfortunately, this methodological caution has been widely ignored.

A variety of research designs can be used to study different questions about System 2, such as the effects of training and intelligence or the efficacy of cues. Dual-task methods are most useful to test hypotheses about the existence of an underlying intuitive judgment that is modified by a corrective intervention of System 2. The test is not whether the judgment will be disrupted by a competing task—such a test would produce too many false positives. The specific prediction is that interference will cause judgments to become more similar to what they would be if System 2 had not had an opportunity to intervene.

The possible corrections in the Linda and Tom W. problems illustrate two possible outcomes of the intervention of System 2. In the case of the conjunction fallacy, the intuitive judgment would be rejected and replaced by an

other conclusion. In the case of base-rate neglect, the intuitive judgment could be adjusted to accommodate a new consideration. Because it is based on salient information and because it comes first, the intuitive impression is likely to serve as an anchor for subsequent adjustments, and the corrective adjustments are therefore likely to be small. Variations on this theme of anchoring on intuition are common in the literature (Epley & Gilovich, 2002; Epstein, 1994; Gilbert, 2002; Griffin & Tversky, 1992; Haidt, 2001; Hammond, 2000; Sloman, 2002; Wilson, 2000; Wilson, Centerbar, & Brekke, 2002). On the other hand, there are reports in which the correction just eliminates the bias of the heuristic judgment (see, e.g., Schwarz, Bless, et al., 1991), as well as occasional findings of over-correction (Oppenheimer, in press). A plausible hypothesis is that adjustments that are based on the explicit identification of a bias are more likely to overcorrect, whereas adjustments based on the identification of an additional relevant considerations are generally insufficient.

Prototype Heuristics

This section introduces a family of prototype heuristics, which share a common mechanism and a remarkably consistent pattern of cognitive illusions, analogous to the effects observed in the Tom W. and in the Linda problems (Kahneman & Frederick, 2002). Prototype heuristics can be roughly described as the substitution of an average for a sum—a process that has been extensively studied by Anderson in other contexts (e.g., Anderson, 1981, pp. 58–70, 1991a, 1991b, 1996). The study of prototype heuristics also illustrates the conditions under which System 2 prevents or reduces judgment biases.

Extensional and Prototype Attributes

The target assessments in several significant tasks of judgment and decision making are *extensional attributes* of categories or sets. The value of an extensional attribute in a set is an aggregate (not necessarily additive) of the values over its extension. Each of the following tasks is illustrated by an example of an extensional attribute and also by the relevant measure of extension. The argument of this section is that the target attributes in these tasks are low in accessibility and are therefore candidates for heuristic judgment.

1. Category prediction (e.g., *the probability that the set of bank tellers contains Linda/the number of bank tellers*);
2. Pricing a quantity of public or private goods (e.g., *the personal dollar value of saving a certain number of birds from drowning in oil ponds/the number of birds*);
3. Global evaluation of a past experience that extended over time (e.g., *the overall aversiveness of a painful medical procedure/the duration of the procedure*); and
4. Assessment of the support that a sample of observations provides for a hypothesis (e.g., *the probability that a specified sample of colored balls has*

been drawn from one urn rather than another/the number of balls).

The logic of extensional attributes involves a general principle of conditional adding, which dictates that each element of the set adds to the overall value an amount that depends on the elements already included. In simple cases, the value is additive: The total length of the set of lines in Figure 3 is just the sum of their separate lengths. In other cases, each positive element of the set increases the aggregate value, but the combination rule is nonadditive (typically, subadditive).⁴

A category or set that is sufficiently homogeneous to have a prototype can also be described by its *prototype attributes*. Where extensional attributes are akin to a sum, prototype attributes are averages. As the display of lines in Figure 3 illustrates, prototype attributes are often highly accessible. This observation is well documented. Whenever people look at, or think about, an ensemble or category that has a prototype, information about the prototype becomes accessible. The classic discussion of basic-level categories included demonstrations of the ease with which features of the prototype come to mind (Rosch & Mervis, 1975). Even earlier, Posner and Keele (1968, 1970) had reported experiments in which observers were exposed on many trials to various distortions of a single shape. The prototype shape was never shown, but observers erroneously believed that it had been presented often. More recently, several studies in social psychology have shown that exposure to the name of a familiar social category increases the accessibility of the traits that are closely associated with its stereotype (see Fiske, 1998).

Because of their high accessibility, the prototype attributes are natural candidates for the role of heuristic attributes. A *prototype heuristic* is the label for the process of substituting an attribute of a prototype for an extensional attribute of its category (Kahneman & Frederick, 2002). The original instance of a prototype heuristic is the use of representativeness in category prediction. The probability of Linda being a bank teller is an extensional variable, but her resemblance to a typical bank teller is a prototype attribute.

Two Tests of Prototype Heuristics

Because extensional and prototypical attributes are governed by characteristically different logical rules, the substitution of a prototype attribute for an extensional attribute entails two testable biases: extension neglect and violations of monotonicity. Tests of the two hypotheses are discussed in turn.

Tests of extension neglect. Doubling the frequencies of all values in a set does not affect prototype attributes because measures of central tendency depend only on relative frequencies. In contrast, the value of an extensional attribute increases monotonically with extension. The hypothesis that judgments of a target attribute are mediated by a prototype heuristic gains support if the judgments are insensitive to variations of extension.

The proposition that extension is neglected in a par-


ticular judgment has the character of a null hypothesis: It is strictly true only if all individuals in the sample are completely insensitive to variations of extension. The hypothesis must be rejected, in a sufficiently large study, if even a small proportion of participants show some sensitivity to extension. The chances of some individuals responding to extension are high a priori because educated respondents are generally aware of the relevance of this variable (Kahneman & Frederick, 2002). Everyone agrees that willingness to pay (WTP) for saving birds should increase with the number of birds saved, that extending a painful medical procedure by an extra period of pain makes it worse, and that evidence from larger samples is more reliable. Complete extension neglect is therefore an unreasonably strict test of prototype heuristics. Nevertheless, this extreme result can be obtained under favorable conditions, as the following examples show.

- The study of Tom W. (see Figure 8) illustrates a pattern of base-rate neglect in categorical prediction. This finding is robust when the task requires a ranking of multiple outcomes (Kahneman & Tversky, 1973). As noted in the preceding section, the sophisticated participants in this experiment were aware of the base rates and were capable of using this knowledge in their predictions—but the thought of doing so apparently occurred to almost none of them. Kahneman and Tversky also documented almost complete neglect of base rates in an experiment (the engineer/lawyer study) in which base rates were explicitly stated. However, the neglect of explicit base-rate information in this design is a fragile finding (see Evans, Handley, Over, & Perham, 2002; Kahneman & Frederick, 2002; Koehler, 1996).
- Participants in a study by Desvousges et al. (1993) indicated their willingness to contribute money to prevent the drowning of migratory birds. The number of birds that would be saved was varied for different subsamples. The estimated amounts that households were willing to pay were \$80, \$78, and \$88, to save 2,000, 20,000, or 200,000 birds, respectively. Frederick and Fischhoff (1998) reviewed numerous other demonstrations of *scope neglect* in studies of WTP for public goods. For example, Kahneman and Knetsch found that survey respondents in Toronto were willing to pay similar amounts to clean up the lakes in a small region of Ontario or to clean up all the lakes in that province (reported by Kahneman, 1986).
- In a study described by Redelmeier and Kahneman

⁴ If the judgment is monotonically related to an additive scale (such as the underlying count of the number of birds), the formal structure is known in the measurement literature as an *extensive structure* (Luce, Krantz, Suppes, & Tversky, 1990, chapter 3). There also may be attributes that lack any underlying additive scale, in which case the structure is known in the literature as a *positive concatenation structure* (Luce et al., 1990, chapter 19, Vol. III, p. 38).

Figure 9

Pain Intensity Reported by Two Colonoscopy Patients



(1996), patients undergoing colonoscopy reported the intensity of pain every 60 seconds during the procedure (see Figure 9) and subsequently provided a global evaluation of the pain they had suffered. The correlation of global evaluations with the duration of the procedure (which ranged from 4 to 66 minutes in that study) was .03. On the other hand global evaluations were correlated ($r = .67$) with an average of the pain reported at two points of the procedure: when pain was at its peak and just before the procedure ended. For example, Patient A in Figure 9 reported a more negative evaluation of the procedure than Patient B. The same pattern of *duration neglect* and *peak/end evaluations* has been observed in other studies (Fredrickson & Kahneman, 1993; see Kahneman, 2000a, 2000b, for a discussion).

In light of the findings discussed in the preceding section, it is useful to consider situations in which people do not neglect extension completely. Extension effects are expected, in the present model, if the individual (a) has information about the extension of the relevant set, (b) is reminded of the relevance of extension, and (c) is able to detect that her intuitive judgment neglects extension. These conditions are least likely to hold—and complete neglect most likely to be observed—when the judge evaluates a single object and when the extension of the set is not explicitly mentioned. At the other extreme, the conditions for a positive effect of extension are all satisfied in psychologists' favorite research design: the within-participant factorial experiment, in which values of extension are crossed with the values of other variables in the design. As noted earlier, this design provides an obvious cue that the experimenter considers every manipulated variable relevant, and it enables participants to ensure that their judgments exhibit sensitivity to all these variables. The factorial

design is therefore especially inappropriate for testing hypotheses about biases of neglect (Kahneman & Frederick, 2002).

Tests of monotonicity. Extensional variables, like sums, obey monotonicity. The sum of a set of positive values is at least as high as the maximum of its subsets. In contrast, the average of a subset can be higher than the average of a set that includes it. Violations of monotonicity are therefore bound to occur when an extensional attribute is judged by a prototype attribute: It is always possible to find cases in which adding elements to a set causes the judgment of the target variable to decrease. This test of prototype heuristics is less demanding than the hypothesis of extension neglect, and violations of monotonicity are compatible with some degree of sensitivity to extension (Ariely & Loewenstein, 2000). Nevertheless, the systematic violation of monotonicity in important tasks of judgment and choice is the strongest source of support for the hypothesis that prototype attributes are being substituted for extensional attributes in these tasks.

- Conjunction errors, which violate monotonicity, have been demonstrated in the Linda problem and in other problems of the same type. The pattern is robust when the judgments are obtained in a between-participants design and when the critical outcomes are embedded in a longer list (Mellers, Hertwig, & Kahneman, 2001; Tversky & Kahneman, 1982, 1983). Tversky and Kahneman (1983) also found that statistically naïve respondents made conjunction errors even in a direct comparison of the critical outcomes. As in the case of extension neglect, however, conjunction errors are less robust in within-participant conditions, especially when the task involves a direct comparison (see Kahneman & Frederick, 2002, for a discussion).
- Hsee (1998) asked participants to price sets of din-

nerware offered in a clearance sale. One of the sets consisted of 24 pieces, all in good condition. The other set included the same 24 pieces, plus 16 additional pieces, of which 7 were in a good condition and 9 were broken. Hsee drew an important distinction between two experimental conditions. When each respondent evaluated only one set of dinnerware (*separate evaluation*), mean WTP was \$33 for the smaller set and \$23 for the larger set ($p < .01$). In contrast, participants who evaluated both sets (*joint evaluation*) were consistently willing to pay more for the larger set. List (2002) observed similar violations of dominance with a different good (sets of baseball cards) in a real market situation.

- Problems of the following kind have been used in several experiments (Griffin & Tversky, 1992; Kahneman & Tversky, 1972).

A sample has been drawn from one of two urns. One urn contains 70% red balls and 30% white balls. The proportions are reversed in the other urn. What is the probability that each of these samples was drawn from the predominantly red urn?

A sample of three red balls and zero white balls (3R, 0W)

A sample of four red balls and three white balls (4R, 3W)

A sample of seven red balls and three white balls (7R, 3W)

- The extensional target variable here is the degree of support for the "red" hypothesis relative to the "white" hypothesis. The normative solution is straightforward: Posterior probability (the target attribute) is determined by an additive combination over sample elements—the difference between the number of red and white balls in the sample. The psychological solution is equally straightforward: The prototype attribute (the heuristic) is an average of support, which corresponds to the proportion of red balls in the sample. Thus, the addition of 4R, 3W to 3R, 0W raises the value of the target attribute but reduces the value of the heuristic attribute. This particular example is fictitious, but the pattern of findings indicates that respondents would derive much more confidence from 3R, 0W than from 7R, 3W (Griffin & Tversky, 1992; Kahneman & Tversky, 1972).
- A randomized clinical experiment was conducted as a follow-up to the colonoscopy study described earlier. For half the patients, the instrument was not immediately removed when the clinical examination ended. Instead, the physician waited for about a minute, leaving the instrument stationary. The experience during the extra period was uncomfortable, but the procedure guaranteed that the colonoscopy never ended in severe pain. Patients reported significantly more favorable global evaluations in this experimental condition than in the control condition

(Redelmeier, Katz, & Kahneman, 2003). Violations of dominance have also been confirmed in choices. Kahneman, Fredrickson, Schreiber, and Redelmeier (1993) exposed participants to two cold-pressor experiences, one with each hand: a short episode (immersion of one hand in 14 °C water for 60 seconds) and a long episode (the short episode plus an additional 30 seconds during which the water was gradually warmed to 15 °C). When participants were later asked which of the two experiences they preferred to repeat, a substantial majority chose the long trial. This pattern of choices is predicted from the peak/end rule of evaluation, which was described earlier. The basic result was replicated with unpleasant sounds of variable loudness and duration (Schreiber & Kahneman, 2000).

The consistency of the results observed in diverse studies of prototype heuristics implies a need for a unified interpretation and challenges interpretations that apply to only a single domain. A number of authors have offered competing interpretations of base-rate neglect (Cosmides & Tooby, 1996; Koehler, 1996), insensitivity to scope in WTP (Kopp, 1992), and duration neglect (Ariely & Loewenstein, 2000). However, each of these interpretations is specific to a particular task and does not carry over to other demonstrations of extension neglect. Similarly, the attempts to describe the conjunction fallacy as a miscommunication between experimenter and respondent (Dulany & Hilton, 1991; Hilton & Slugoski, 2001) do not explain analogous violations of monotonicity in the cold-pressor experiment and in the pricing of private goods. In contrast, the account offered here (and developed in greater detail by Kahneman & Frederick, 2002) is equally applicable to diverse tasks that require an assessment of an extensional target attribute. Future discussions of the separate phenomena should take account of their generality across domains.

The findings obtained in choices and joint evaluations confirm the existence of two distinct ways of choosing, which were already identified in prospect theory (Kahneman & Tversky, 1979). In the nonanalytic procedure that I have called *choosing by liking* (Kahneman, 1994), the individual considers the global evaluation of the two options separately and selects the option that has the higher global value without detailed comparison of the alternatives. Choice by global value is the basic mechanism assumed in prospect theory. However, prospect theory also introduces the idea that if the decision maker detects that one option dominates the other, she will choose the dominant option without consulting their separate valuations. The same mechanisms apply to problems of judgment, such as the case of Linda, where some statistically sophisticated individuals detect that one of the sets includes the other and respond accordingly, ignoring representativeness. In Hsee's (1998) dinnerware study, respondents chose by liking in separate evaluation and chose by dominance in joint evaluation.

Joint evaluation is not sufficient to guarantee choice by dominance; it is also necessary for the decision makers

to realize explicitly that one of the options is strictly better than the other. This requirement was not satisfied in the cold-pressor experiment. Although the participants were exposed to both experiences (joint evaluation), they did not notice that the long episode contained all the pain of the short one and then some extra pain. Most respondents would have made a different choice if they had understood the structure of the options.

The cases that have been discussed are only illustrations, not a comprehensive list of prototype heuristics. For example, the same form of nonextensional thinking explains why the median estimate of the annual number of murders in Detroit is twice as high as the estimate of the number of murders in Michigan (Kahneman & Frederick, 2002). It also explains why professional forecasters assigned a higher probability to "an earthquake in California causing a flood in which more than 1,000 people will drown" than to "a flood somewhere in the United States in which more than 1,000 people will drown" (Tversky & Kahneman, 1983).

The normative logic of belief and choice is extensional, and it requires appropriate valuation of extensional attributes, which include both probability and utility. The examples that were discussed in this section demonstrate pervasive violations of extensional logic in the intuitive evaluation of both evidence and outcomes. The substitution of prototype attributes for extensional attributes appears to be a general characteristic of System 1, which is incompatible with both Bayesian beliefs and utility maximization.

Conclusions

The starting point of the present analysis was the observation that complex judgments and preferences are called intuitive in everyday language if they come to mind quickly and effortlessly, like percepts. Other basic observations were that judgments and intentions are normally intuitive in this sense and that they can be modified or overridden in a more deliberate mode of operation. The labels *System 1* and *System 2* were associated with these two modes of cognitive functioning.

The preceding sections elaborated a single generic proposition: Highly accessible impressions produced by System 1 control judgments and preferences, unless modified or overridden by the deliberate operations of System 2. This template sets an agenda for research: To understand judgment and choice, we must study the determinants of high accessibility, the conditions under which System 2 overrides or corrects System 1, and the rules of these corrective operations. Much is known about each of the three questions.

First, consider the ways in which the concept of accessibility was used here. Framing effects were attributed to the fact that alternative formulations of the same situation make different aspects of it accessible. The core idea of prospect theory, that the normal carriers of utility are gains and losses, invoked the general principle that changes are relatively more accessible than absolute values. Judgment heuristics were explained as the substitution of a

highly accessible heuristic attribute for a less accessible target attribute. The correction of intuitive judgments was attributed to the accessibility of competing considerations and to the accessibility of metacognitive awareness of bias. Finally, the proposition that averages are more accessible than sums unified the analysis of prototype heuristics. A recurrent theme was that different aspects of problems are made accessible in between-participants and in within-participant experiments and more specifically in separate and joint evaluations of stimuli. In all these contexts, the discussion appealed to rules of accessibility that are independently plausible and sometimes quite obvious.

As was noted earlier, the status of accessibility factors in psychological theorizing is, in principle, similar to the status of perceptual grouping factors. In both cases, there is no general theory, only a list of powerful empirical generalizations that provide an adequate basis for experimental predictions and for models of higher level phenomena. Unlike gestalt principles, which were catalogued a long time ago, a comprehensive list of the factors that influence accessibility is yet to be drawn. The list will be long, but many of its elements are already known. For example, it is safe to assume that similarity is more accessible than probability, that changes are more accessible than absolute values, that averages are more accessible than sums, and that the accessibility of a rule of logic or statistics can be temporarily increased by a reminder. Furthermore, each of these assumptions can be verified independently by multiple operations, including measurements of reaction time, susceptibility to interference by secondary tasks, and asymmetric priming effects. Assumptions about accessibility are incompletely theorized, but they need not be vague, and they can do genuine explanatory work. The claim "X came to mind because it was accessible under the circumstances of the moment" sounds circular, but it is not.

The discussion of judgment heuristics was restricted to the differential accessibility of attributes (dimensions) on which judgment objects vary, such as length or price, similarity and probability (Kahneman & Frederick, 2002). A similar analysis could be applied to the accessibility of particular values of attributes, such as "six feet" or "two dollars." Highly accessible values are generally overweighted, and when considered as possible answers to a question, they become potent anchors (Chapman & Johnson, 2002; Epley & Gilovich, 2002; Strack & Mussweiler, 1997). These effects of salience and anchoring play a central role in treatments of judgment and choice. Indeed, anchoring effects are among the most robust phenomena of judgment, and overweighting of salient values is likely to be the mechanism that explains why low-probability events sometimes loom large in decision making. The analysis of accessibility could readily be extended to deal with these observations.

The claim that cognitive illusions occur unless they are prevented by System 2 also sounds circular but is not. Circular inferences are avoidable because the role of System 2 can be independently verified in several ways. For example, the assumption that System 2 is vulnerable to interference by competing activities suggests that manifes-

tations of intuitive thought that are normally inhibited will be expressed when people are placed under cognitive load. Another testable hypothesis is that intuitive judgments that are suppressed by System 2 still have detectable effects, for example, in priming subsequent responses.

Principles of accessibility determine the relative power of the cues to which the monitoring functions of System 2 respond. For example, it is known that differences between objects of choice or judgment are more salient in joint than in separate evaluation, and it is also known that any variable that is manipulated in a factorial design will attract some attention. Other cues can be found in the wording of problems and in the context of previous tasks. Many apparent inconsistencies in the literature on judgment heuristics are easily resolved within this framework (Kahneman & Frederick, 2002). The observation that a judgment bias appears in some situations but not in others usually provides information about the accessibility factors that trigger corrective operations. As already noted, the attribution of the variability of intuitive judgments to System 2 is a source of testable hypotheses. It suggests, for example, that intelligence will be correlated with susceptibility to biases only in problems that provide relatively weak cues to the correct solution. In the absence of cues, there is no opportunity for intelligence or sophistication to manifest itself. At the other extreme, when cues are abundant, even the moderately intelligent will find them (Kahneman, 2000c; Stanovich & West, 1999, 2002).

The model suggests five ways in which a judgment or choice may be made:

1. An intuitive judgment or intention is initiated, and
 - (a) Endorsed by System 2;
 - (b) Adjusted (insufficiently) for other features that are recognized as relevant;
 - (c) Corrected (sometimes overcorrected) for an explicitly recognized bias; or
 - (d) Identified as violating a subjectively valid rule and blocked from overt expression.
2. No intuitive response comes to mind, and the judgment is computed by System 2.

There is, of course, no way to ascertain precisely the relative frequencies of these outcomes, but casual observation suggests that Cases 1(a) and 1(b) are the most common and that Case 1(d) is very rare. This ordering reflects two major hypotheses about the role of intuition in judgment and choice. The first is that most behavior is intuitive, skilled, unproblematic, and successful (Klein, 1998). The second is that behavior is likely to be anchored in intuitive impressions and intentions even when it is not completely dominated by them. An essay with a related message (Haidt, 2001) suggested the image of the intuitive dog wagging the rational tail.

The analysis of intuitive thinking and choice that has been presented here provides a framework that highlights commonalities between lines of research that are usually studied separately. In particular, the psychology of judgment and the psychology of choice share their basic principles and differ mainly in content. At a more specific level,

prototype heuristics solve structurally similar problems in diverse domains, where they yield closely similar patterns of results. Furthermore, the psychological principles that have been invoked are not specific to the domain of judgment/decision making. The analogy between intuition and perception has been especially fruitful in identifying the ways in which intuitive thought differs from deliberate reasoning, and the notions of accessibility and dual-process analyses play a fundamental role in other domains of social and cognitive psychology.

A general framework such as the one offered here is not a substitute for domain-specific concepts and theories. For one thing, general frameworks and specific models make different ideas accessible. Novel ideas and compelling examples are perhaps more likely to arise from thinking about problems at a lower level of abstraction and generality. However, broad concepts such as accessibility, attribute substitution, corrective operations, and prototype heuristics can be useful if they guide a principled search for analogies across domains, help identify common processes, and prevent overly narrow interpretations of findings.

REFERENCES

- Agnoli, F. (1991). Development of judgmental heuristics and logical reasoning: Training counteracts the representativeness heuristic. *Cognitive Development*, 6, 195-217.
- Agnoli, F., & Krantz, D. H. (1989). Suppressing natural heuristics by formal instruction: The case of the conjunction fallacy. *Cognitive Psychology*, 21, 515-550.
- Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.
- Anderson, N. H. (1991a). *Contributions to information integration theory: Vol. I. Cognition*. Hillsdale, NJ: Erlbaum.
- Anderson, N. H. (1991b). *Contributions to information integration theory: Vol. II. Social*. Hillsdale, NJ: Erlbaum.
- Anderson, N. H. (1996). *A functional theory of cognition*. Hillsdale, NJ: Erlbaum.
- Ariely, D. (2001). Seeing sets: Representation by statistical properties. *Psychological Science*, 12, 157-162.
- Ariely, D., & Loewenstein, G. (2000). When does duration matter in judgment and decision making? *Journal of Experimental Psychology: General*, 129, 524-529.
- Arrow, K. J. (1982). Risk perception in psychology and economics. *Economic Inquiry*, 20, 1-9.
- Bargh, J. A. (1997). The automaticity of everyday life. In R. S. Wyer, Jr. (Ed.), *The automaticity of everyday life: Advances in social cognition* (Vol. 10, pp. 1-61). Mahwah, NJ: Erlbaum.
- Bar-Hillel, M., & Neter, E. (2002). How alike is it versus how likely is it: A disjunction fallacy in probability judgments. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (pp. 82-97). New York: Cambridge University Press.
- Benartzi, S., & Thaler, R. H. (1999). Risk aversion or myopia? Choices in repeated gambles and retirement investments. *Management Science*, 47, 364-381.
- Bernoulli, D. (1954). Exposition of a new theory on the measurement of risk. *Econometrica*, 22, 23-36. (Original work published 1738)
- Bless, H., Clore, G. L., Schwarz, N., Golisano, V., Rabe, C., & Wolk, M. (1996). Mood and the use of scripts: Does a happy mood really lead to mindlessness? *Journal of Personality and Social Psychology*, 71, 665-679.
- Bodenhausen, G. V. (1990). Stereotypes as judgmental heuristics: Evidence of circadian variations in discrimination. *Psychological Science*, 1, 319-322.
- Brenner, L. A., Koehler, D. J., & Rottenstreich, Y. (2002). Remarks on support theory: Recent advances and future directions. In T. Gilovich,

- effect, loss aversion, and status quo bias. *Journal of Economic Perspectives*, 5, 193–206.
- Kahneman, D., & Lovallo, D. (1993). Timid choices and bold forecasts: A cognitive perspective on risk taking. *Management Science*, 39, 17–31.
- Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, 93, 136–153.
- Kahneman, D., & Ritov, I. (1994). Determinants of stated willingness to pay for public goods: A study in the headline method. *Journal of Risk and Uncertainty*, 9, 5–38.
- Kahneman, D., Ritov, I., & Schkade, D. (1999). Economic preferences or attitude expressions? An analysis of dollar responses to public issues. *Journal of Risk and Uncertainty*, 19, 220–242.
- Kahneman, D., Schkade, D. A., & Sunstein, C. R. (1998). Shared outrage and erratic awards: The psychology of punitive damages. *Journal of Risk and Uncertainty*, 16, 49–86.
- Kahneman, D., Slovic, P., & Tversky, A. (Eds.). (1982). *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.
- Kahneman, D., & Tversky, A. (1972). Subjective probability: A judgment of representativeness. *Cognitive Psychology*, 3, 430–454.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237–251.
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decisions under risk. *Econometrica*, 47, 263–291.
- Kahneman, D., & Tversky, A. (1982). On the study of statistical intuitions. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 493–508). New York: Cambridge University Press.
- Kahneman, D., & Tversky, A. (Eds.). (2000). *Choices, values, and frames*. New York: Cambridge University Press.
- Kahneman, D., & Varey, C. A. (1990). Propensities and counterfactuals: The loser that almost won. *Journal of Personality and Social Psychology*, 59, 1101–1110.
- Kahneman, D., Wakker, P. P., & Sarin, R. (1997). Back to Bentham? Explorations of experienced utility. *Quarterly Journal of Economics*, 112, 375–405.
- Keren, G., & Wagenaar, W. A. (1987). Violations of utility theory in unique and repeated gambles. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13, 387–391.
- Klein, G. (1998). *Sources of power: How people make decisions*. Cambridge, MA: MIT Press.
- Klein, G. (2003). *Intuition at work: Why developing your gut instincts will make you better at what you do*. New York: Doubleday.
- Koehler, D. J. (1996). A strength model of probability judgments for tournaments. *Organizational Behavior and Human Decision Processes*, 66, 16–21.
- Kopp, R. (1992). Why existence value should be used in cost-benefit analysis. *Journal of Policy Analysis and Management*, 11, 123–130.
- LeBoeuf, R. A., & Shafir, E. (2003). Deep thoughts and shallow frames: On the susceptibility to framing effects. *Journal of Behavioral Decision Making*, 16, 77–92.
- LeDoux, J. E. (2000). Emotion circuits in the brain. *Annual Review of Neuroscience*, 23, 155–184.
- List, J. (2002). Preference reversals of a different kind: The more is less phenomenon. *American Economic Review*, 92, 1636–1643.
- Loewenstein, G. (1996). Out of control: Visceral influences on behavior. *Organizational Behavior and Human Decision Processes*, 65, 272–292.
- Loewenstein, G., Weber, E. U., Hsee, C. K., & Welch, N. (2001). Risk as feelings. *Psychological Bulletin*, 127, 267–286.
- Luce, R. D., Krantz, D. H., Suppes, P., & Tversky, A. (1990). *Foundations of measurement: Vol. 3. Representation, axiomatization, and invariance*. San Diego, CA: Academic Press.
- McNeil, B. J., Pauker, S. G., Sox, H. C., & Tversky, A. (1982). On the elicitation of preferences for alternative therapies. *New England Journal of Medicine*, 306, 1259–1262.
- Mellers, B. (2000). Choice and the relative pleasure of consequences. *Psychological Bulletin*, 126, 910–924.
- Mellers, B., Hertwig, R., & Kahneman, D. (2001). Do frequency representations eliminate conjunction effects? An exercise in adversarial collaboration. *Psychological Science*, 12, 269–275.
- Michotte, A. (1963). *The perception of causality* (T. R. Miles & E. Miles, Trans.). New York: Basic Books.
- Myers, D. G. (2002). *Intuition: Its powers and perils*. New Haven, CT: Yale University Press.
- Nisbett, R. E., Krantz, D. H., Jepson, C., & Kunda, Z. (1983). The use of statistical heuristics in everyday inductive reasoning. *Psychological Review*, 90, 339–363.
- Oppenheimer, D. M. (in press). Spontaneous discounting of availability in frequency estimations. *Psychological Science*.
- Palmer, S. E. (1999). *Vision science: Photons to phenomenology*. Cambridge, MA: MIT Press.
- Pashler, H. E. (1998). *The psychology of attention*. Cambridge, MA: MIT Press.
- Posner, M. I., & Keele, S. W. (1968). On the genesis of abstract ideas. *Journal of Experimental Psychology*, 77, 353–363.
- Posner, M. I., & Keele, S. W. (1970). Retention of abstract ideas. *Journal of Experimental Psychology*, 83, 304–308.
- Poulton, E. C. (1975). Range effects in experiments with people. *American Journal of Psychology*, 77, 353–363.
- Redelmeier, D., & Kahneman, D. (1996). Patients' memories of painful medical treatments: Real-time and retrospective evaluations of two minimally invasive procedures. *Pain*, 66, 3–8.
- Redelmeier, D. A., Katz, J., & Kahneman, D. (2003). Memories of colonoscopy: A randomized trial. *Pain*, 104, 187–194.
- Rosch, E., & Mervis, C. B. (1975). Family resemblances: Studies in the internal structure of categories. *Cognitive Psychology*, 7, 573–605.
- Samuelson, W., & Zeckhauser, R. (1988). Status quo bias in decision making. *Journal of Risk and Uncertainty*, 1, 7–59.
- Schreiber, C. A., & Kahneman, D. (2000). Determinants of the remembered utility of aversive sounds. *Journal of Experimental Psychology: General*, 129, 27–42.
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatke, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology*, 61, 195–202.
- Schwarz, N., & Clore, G. L. (1983). Mood, misattribution, and judgments of well-being: Informative and directive functions of affective states. *Journal of Personality and Social Psychology*, 45, 513–523.
- Schwarz, N., & Strack, F. (1999). Reports of subjective well-being: Judgmental processes and their methodological implications. In D. Kahneman, E. Diener, & N. Schwarz (Eds.), *Well-being: The foundations of hedonistic psychology* (pp. 61–84). New York: Russell Sage Foundation.
- Schwarz, N., Strack, F., & Mai, H. P. (1991). Assimilation and contrast effects in part-whole question sequences: A conversational logic analysis. *Public Opinion Quarterly*, 55, 3–23.
- Schwarz, N., & Vaughn, L. A. (2002). The availability heuristic revisited: Ease of recall and content of recall as distinct sources of information. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (pp. 103–119). New York: Cambridge University Press.
- Shafir, E. (1993). Choosing versus rejecting: Why some options are both better and worse than others. *Memory and Cognition*, 21, 546–556.
- Shafir, E., & LeBoeuf, R. A. (2002). Rationality. *Annual Review of Psychology*, 53, 419–517.
- Simon, H. A. (1955). A behavioral model of rational choice. *Quarterly Journal of Economics*, 69, 99–118.
- Simon, H. A. (1979). Rational decision-making in business organizations. *American Economic Review*, 69, 493–513.
- Simon, H. A., & Chase, W. G. (1973). Skill in chess. *American Scientist*, 61, 394–403.
- Simon, H. A., & Hayes, J. R. (1976). Understanding process: Problem isomorphs. *Cognitive Psychology*, 8, 165–190.
- Slooman, S. A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin*, 119, 3–22.
- Slooman, S. A. (2002). Two systems of reasoning. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (pp. 379–396). New York: Cambridge University Press.
- Slovic, P., Finucane, M., Peters, E., & MacGregor, D. G. (2002). The affect heuristic. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (pp. 397–420). New York: Cambridge University Press.

- Smith, S. M., & Levin, I. P. (1996). Need for cognition and choice framing effects. *Journal of Behavioral Decision Making*, 9, 283-290.
- Stanovich, K. E. (1999). *Who is rational? Studies of individual differences in reasoning*. Mahwah, NJ: Erlbaum.
- Stanovich, K. E., & West, R. F. (1999). Discrepancies between normative and descriptive models of decision making and the understanding/acceptance principle. *Cognitive Psychology*, 38, 349-385.
- Stanovich, K. E., & West, R. F. (2000). Individual differences in reasoning: Implications for the rationality debate. *Behavioral and Brain Sciences*, 23, 645-665.
- Stanovich, K. E., & West, R. F. (2002). Individual differences in reasoning: Implications for the rationality debate. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (pp. 421-440). New York: Cambridge University Press.
- Strack, F., Martin, L. L., & Schwarz, N. (1988). Priming and communication: The social determinants of information use in judgments of life-satisfaction. *European Journal of Social Psychology*, 18, 429-442.
- Strack, F., & Mussweiler, T. (1997). Explaining the enigmatic anchoring effect: Mechanisms of selective accessibility. *Journal of Personality and Social Psychology*, 73, 437-446.
- Swalm, R. O. (1966). Utility theory: Insights into risk taking. *Harvard Business Review*, 44, 123-136.
- Thaler, R. H. (1980). Toward a positive theory of consumer choice. *Journal of Economic Behavior and Organization*, 1, 39-60.
- Thaler, R. H. (1985). Mental accounting and consumer choice. *Marketing Science*, 4, 199-214.
- Thaler, R. H. (1999). Mental accounting matters. *Journal of Behavioral Decision Making*, 12, 241-268.
- Tulving, E., & Pearlstone, Z. (1966). Availability versus accessibility of information in memory for words. *Journal of Verbal Learning and Verbal Behavior*, 5, 381-391.
- Tversky, A., & Kahneman, D. (1971). Belief in the law of small numbers. *Psychological Bulletin*, 76, 105-110.
- Tversky, A., & Kahneman, D. (1973). Availability: A heuristic for judging frequency and probability. *Cognitive Psychology*, 5, 207-232.
- Tversky, A., & Kahneman, D. (1974, September 27). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124-1131.
- Tversky, A., & Kahneman, D. (1981, January 30). The framing of decisions and the psychology of choice. *Science*, 211, 453-458.
- Tversky, A., & Kahneman, D. (1982). Evidential impact of base rates. In D. Kahneman, P. Slovic, & A. Tversky (Eds.), *Judgment under uncertainty: Heuristics and biases* (pp. 153-160). New York: Cambridge University Press.
- Tversky, A., & Kahneman, D. (1983). Extensional vs. intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293-315.
- Tversky, A., & Kahneman, D. (1986). Rational choice and the framing of decisions. *Journal of Business*, 59, S251-S278.
- Tversky, A., & Kahneman, D. (1991). Loss aversion in riskless choice: A reference-dependent model. *Quarterly Journal of Economics*, 106, 1039-1061.
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, 5, 297-323.
- Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review*, 101, 547-567.
- Tversky, A., & Redelmeier, D. A. (1992). On the framing of multiple prospects. *Psychological Science*, 3, 191-193.
- Wilson, T. D. (2000). *Strangers to ourselves: Discovering the adaptive unconscious*. Cambridge, MA: Harvard University Press.
- Wilson, T. D., Centerbar, D. B., & Brekke, N. (2002). Mental contamination and the debiasing problem. In T. Gilovich, D. Griffin, & D. Kahneman (Eds.), *Heuristics and biases* (pp. 185-200). New York: Cambridge University Press.
- Wilson, T. D., & Schooler, J. W. (1991). Thinking too much: Introspection can reduce the quality of preferences and decisions. *Journal of Personality and Social Psychology*, 60, 181-192.
- Witreich, W. J. (1961). The Honi phenomenon: A case of selective perceptual distortion. In F. P. Kilpatrick (Ed.), *Explorations in transactional psychology* (pp. 188-202). New York: New York University Press.
- Zajonc, R. B. (1980). Feeling and thinking: Preferences need no inferences. *American Psychologist*, 35, 151-175.
- Zajonc, R. B. (1998). Emotions. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology* (4th ed., Vol. 1, pp. 591-632). New York: Oxford University Press.
- Zukier, H., & Pepitone, A. (1984). Social roles and strategies in prediction: Some determinants in the use of base-rate information. *Journal of Personality and Social Psychology*, 47, 349-360.