

ramifications of Support Theory's insight that judgments are not based on events themselves, but on descriptions of events. The examination of complementary perspectives on heuristics ends with four chapters on specific heuristics beyond those originally proposed by Kahneman and Tversky, and one chapter that examines the implications of considering alternative metaphors of the human judge.

The book concludes with ten chapters on various applications of the heuristics and biases approach to judgment. Four of these deal with judgments made by "the average person" in various aspects of everyday life and six are concerned with the judgments rendered by experts in a number of applied domains. These chapters are significant because they illustrate that the processes of judgment revealed by psychological research are not restricted to the psychological laboratory or to unfamiliar and unimportant tasks. These heuristics – and the biases that are associated with them – have implications for some of the most consequential judgments that life requires people to make.

On a procedural front, we should note that the source of each preexisting piece is indicated by a footnote on the opening page. All preexisting pieces have been edited to some degree. Deletions from the original are indicated by ellipses (...). In nearly all cases, such deletions (and the concomitant renumbering of tables, figures, etc.) constitute the only changes from the original. Exceptions are noted in the opening footnote of the pertinent chapters. Readers interested in the full statement of any author are encouraged to consult the original work. Note also that all references are contained in an overall reference list at the back of the book, and that references to other chapters in the book are indicated by chapter number.

Our work in preparing this book was supported by NSF grants 9809262 and 0091416 to Cornell University, a Social Sciences and Humanities Research Council of Canada research grant to the University of British Columbia, and NSF grant 2556558 to Princeton University.

We wish to thank Zachary Dorsey, Richard Gonzalez, Michael Griffin, and Phil Laughlin for their help in the preparation of this book.

Thomas Gilovich
Dale Griffin
Daniel Kahneman

Introduction – Heuristics and Biases: Then and Now

Thomas Gilovich and Dale Griffin

In the late 1960s and early 1970s, a series of papers by Amos Tversky and Daniel Kahneman revolutionized academic research on human judgment. The central idea of the "heuristics and biases" program – that judgment under uncertainty often rests on a limited number of simplifying heuristics rather than extensive algorithmic processing – soon spread beyond academic psychology, affecting theory and research across a range of disciplines including economics, law, medicine, and political science. The message was revolutionary in that it simultaneously questioned the descriptive adequacy of ideal models of judgment and offered a cognitive alternative that explained human error without invoking motivated irrationality. The initial papers and a variety of related work were collected in a 1982 volume, *Judgment under Uncertainty: Heuristics and Biases* (Kahneman, Slovic, & Tversky, 1982). In the time since, research in the heuristics and biases tradition has prospered on a number of fronts, each represented by a section of the current volume. In this opening chapter, we wish to put the heuristics and biases approach in historical context and discuss some key issues that have been raised since the 1982 book appeared.

HISTORICAL OVERVIEW

Any discussion of the modern history of research on everyday judgment must take note of the large shadow cast by the classical model of rational choice. The model has been applied most vigorously in the discipline of economics, but its considerable influence can be felt in all the behavioral and social sciences and in related policy fields such as law and medicine. According to this model, the "rational actor" (i.e., the typical person) chooses what options to pursue by assessing the probability of each possible outcome, discerning the utility to be derived from each, and combining these two assessments. The option pursued is the one that offers the optimal combination of probability and utility.

Calculations of probability and multiattribute utility can be rather formidable judgments to make, but the theory of rational choice assumes that people make them and make them well. Proponents of the theory do not insist that people never make mistakes in these calculations; but they do insist that the

mistakes are unsystematic. The model assumes, for example, that the rational actor will follow the elementary rules of probability when calculating, say, the likelihood of a given candidate winning an election or the odds of surviving a surgical intervention.

But is the average person as attuned to the axioms of formal rationality as this stance demands? Much of the modern history of judgment research can be summarized as follows. First, evidence is collected indicating that people's assessments of likelihood and risk do not conform to the laws of probability. Second, an argument ensues about the significance of these demonstrations between proponents of human rationality and those responsible for the empirical demonstrations. Three early contributions to this debate – one empirical, one methodological, and one theoretical – have been especially influential.

The empirical contribution was provided by Paul Meehl (1954), who compiled evidence comparing expert clinical prediction with actuarial methods and found that the actuarial methods, or formulas, almost always did better. His research also uncovered a sharp discrepancy between clinicians' assessments of their performance and their actual record of success (see Chapter 40 by Dawes, Faust, & Meehl for a modern summary of this literature). The juxtaposition of modest performance and robust confidence inspired research on faulty processes of reasoning that yield compelling but mistaken inferences.

Ward Edwards made a key methodological contribution by introducing Bayesian analyses to psychology, thus providing a normative standard with which everyday judgments could be compared (Edwards, Lindman, & Savage, 1963). From Edwards' own research (Edwards, 1968) and much that followed, it was clear that intuitive judgments of likelihood did not exactly correspond with this "ideal" normative standard. This led, in turn, to an interest in the causes of suboptimal performance and strategies for improvement.

The most significant theoretical development in this field was Herbert Simon's contention that the "full" rationality implied by the rational choice model was an unrealistic standard for human judgment. He proposed a more limited criterion for actual performance, famously dubbed *bounded rationality*, that acknowledged the inherent processing limitations of the human mind. People reason and choose rationally, but only within the constraints imposed by their limited search and computational capacities. Simon (1957) also discussed the simplifying heuristics that people could employ to cope effectively with these limitations. Note that Simon did not reject the normative appeal of the full-information rational models, referring to them as "jewels of intellectual accomplishment" (Simon, 1983). (Two of the present contributions are strongly influenced by the Simonian perspective on heuristics: Frederick, Chapter 30; and Gigerenzer, Czerlinski, & Martignon, Chapter 31).

The Heuristics and Biases Approach

Inspired by the examples of biased real-world judgments of the sort identified by Meehl and his peers, and guided by the clear normative theory explicated by

Edwards and others, Kahneman and Tversky developed their own perspective on bounded rationality. Although acknowledging the role of task complexity and limited processing capacity in erroneous judgment, Kahneman and Tversky were convinced that the processes of intuitive judgment were not merely simpler than rational models demanded, but were categorically different in kind. Kahneman and Tversky described three general-purpose heuristics – availability, representativeness, and anchoring and adjustment – that underlie many intuitive judgments under uncertainty. These heuristics, it was suggested, were simple and efficient because they piggybacked on basic computations that the mind had evolved to make. Thus, when asked to evaluate the relative frequency of cocaine use in Hollywood actors, one may assess how easy it is to retrieve examples of celebrity drug-users – the availability heuristic piggybacks on highly efficient memory retrieval processes. When evaluating the likelihood that a given comic actor is a cocaine user, one may assess the similarity between that actor and the prototypical cocaine user (the representativeness heuristic piggybacks on automatic pattern-matching processes). Either question may also be answered by starting with a salient initial value (say, 50%) and adjusting downward to reach a final answer (the anchoring and adjustment heuristic, whose underlying mechanisms are debated in Chapters 6 and 7).

In the early experiments that defined this work, each heuristic was associated with a set of *biases*: departures from the normative rational theory that served as markers or signatures of the underlying heuristics. Use of the availability heuristic, for example, leads to error whenever memory retrieval is a biased cue to actual frequency because of an individual's tendency to seek out and remember dramatic cases or because of the broader world's tendency to call attention to examples of a particular (restricted) type. Some of these biases were defined as deviations from some "true" or objective value, but most by violations of basic laws of probability. (For elegant examples, see Chapter 1 by Tversky and Kahneman).

Several aspects of this program are important to note at the outset because they set the stage for a discussion of the criticisms it aroused. First, although the heuristics are distinguished from normative reasoning processes by patterns of biased judgments, the heuristics themselves are sensible estimation procedures that are by no measure "irrational." Second, although heuristics yield "quick and dirty" solutions, they draw on underlying processes (e.g., feature matching, memory retrieval) that are highly sophisticated. Finally, note that these heuristic processes are not exceptional responses to problems of excessive complexity or an overload of information, but normal intuitive responses to even the simplest questions about likelihood, frequency, and prediction.

The Positive and Negative Agendas. As the preceding discussion implies, Kahneman and Tversky distinguished between two messages or agendas for the heuristics and biases program, one "positive" and one "negative." The positive agenda is to elucidate the processes through which people make a variety of important and difficult real world judgments. Is a corporation's explosive

4 Thomas Gilovich and Dale Griffin

growth likely to continue? Is a coup more likely in Ecuador or Indonesia? What is a reasonable estimate of next year's GNP? Thus, representativeness, availability, and anchoring and adjustment were proposed as a set of highly efficient mental shortcuts that provide subjectively compelling and often quite serviceable solutions to such judgmental problems.

But, the solutions were just that – serviceable, not exact or perfectly accurate. Thus the second, negative, agenda of the heuristics and biases program was to specify the conditions under which intuitive judgments were likely to depart from the rules of probability. When, in other words, are everyday judgments likely to be biased? Kahneman and Tversky's experience teaching statistics and their observations of predictions made in applied settings led them to conclude that people often fail to anticipate regression to the mean, fail to give adequate weight to sample size in assessing the import of evidence, and fail to take full advantage of base rates when making predictions. Their three (now familiar) heuristics were offered as an explanation of the when and why of such errors. Thus the two agendas blend together: Identifying particular biases is important in its own right, but doing so also sheds light on the underlying processes of judgment. (Kahneman & Tversky, 1982b, also offered positive and negative approaches to judgment errors, a perspective that is taken up by Kahneman & Frederick, Chapter 2.)

Automatic or Deliberate? There is another dichotomous aspect of the heuristics and biases approach that warrants discussion. Heuristics have often been described as something akin to strategies that people use deliberately in order to simplify judgmental tasks that would otherwise be too difficult for the typical human mind to solve. This use of the term fits with the "cognitive miser" metaphor that proved popular in the field of social cognition (Fiske & Taylor, 1991). The metaphor suggests, perhaps unfortunately and unwisely, that the biases documented in the heuristics and biases tradition are the product of lazy and inattentive minds. The implication is unfortunate and potentially misleading because the biases identified in this tradition have not been appreciably reduced by incentives for participants to sit straight, pay attention, and devote their full cognitive resources to the task (Camerer & Hogarth, 1999; Grether & Plott, 1979; Wilson, Houston, Etling, & Brekke, 1996; see Chapters 37 through 42 for descriptions of real-world judgments characterized by a high level of domain-specific expertise and motivation that nonetheless fit the patterns described by the heuristics and biases program). After reviewing 74 studies, Camerer and Hogarth (1999) concluded that incentives can reduce self-presentation effects, increase attention and effort, and reduce thoughtless responding, but noted that "no replicated study has made rationality violations disappear purely by raising incentives" (p. 7).

Imperviousness to incentives is just what one would expect from considering the other half of the dichotomy, or the other way that heuristics have been described. In particular, Tversky and Kahneman (1983; see Chapter 1) tied heuristics to "natural assessments" elicited by the task at hand that can influence

judgment without being used deliberately or strategically. When deciding whether an elegantly-dressed lawyer is more likely to be a public defender or a member of a large corporate firm, for example, one cannot help computing the similarity between the individual and the prototype of each professional niche. This assessment then informs the judgment of likelihood in the absence of deliberative intent.

It seems to us that both uses of the term are valid and have their place. When deciding whether there are more coups in Ecuador or Indonesia, for example, one automatically searches for known instances of each (availability). Yet one can also deliberately recruit such instances and use the ease with which they come to mind as an explicit strategy – as when deciding to bet on one team over another after explicitly considering the number of star players on each. Similarly, existing research on anchoring makes it clear that many anchoring effects occur in the absence of any explicit adjustment (Mussweiler & Strack, 1999; see Chapman & Johnson, Chapter 6). Often people's estimates are automatically contaminated by previously mentioned values. Sometimes, however, the anchoring and adjustment heuristic is deliberately employed. If asked when George Washington was first elected president, most people do not know the answer; but they do know it was after 1776 and they adjust from that year (Epley & Gilovich, Chapter 7). Anchoring and adjustment is thus sometimes used as an explicit strategy of judgment.

For reasons that have to do with what was going on elsewhere in psychology, the "cognitive miser" view of heuristics took hold more pervasively than the "natural assessments" view, a result that still holds true today. With the rise of "two system" models of reasoning, however (described in Chapter 2 and Chapters 22 through 24), we predict this will change. The two systems view is consistent with the idea of rapid, automatic assessments that may or may not be overridden by deliberate processes, and the emergence of such a perspective should provide a boost to this relatively neglected statement of how heuristics work. Indeed, one of the objectives of this book is to reassert the natural assessments view of heuristics, a stance laid out most forcefully in the opening chapter by Tversky and Kahneman, and discussed extensively in Chapter 2 by Kahneman and Frederick.

Why The Heuristics and Biases Program Had (and Has) Such Influence and Appeal

The impact of any idea is a product of the quality of the idea itself and the intellectual zeitgeist at the time it is offered. Successful ideas must not only be good, but timely – even lucky. So it has been with the heuristics and biases approach to judgment. The popularity and impact of the approach were enhanced by elements of the prevailing atmosphere at the time it was launched, several of which still hold true today. Most important, perhaps, is the very strength and resilience of the rational choice model that motivated much of the heuristics and biases research. Although the model is most entrenched in the field of

economics, it has had a profound impact on theoretical development in sociology, political science, law, and psychology as well. The very reach of the rational choice model thus opens up terrain for any systematic critique that offers an alternative perspective. Wherever the rational choice model shall go, in other words, the heuristics and biases program – or something much like it – must follow. And follow it did, as the heuristics and biases program has reshaped both explicit and implicit assumptions about human thought in all of these areas and a few more.

Models of spending and investment behavior have been particularly influenced by the heuristics and biases program, thanks partly to the deft translations offered by economist Richard Thaler (see DeBondt and Thaler, Chapter 38). Thaler's work is an example of how the heuristics and biases program has become a "full-circle" paradigm: insights that were sparked by observations in the classroom, battlefield, and conference room, then sharpened and tested in the experimental laboratory, are ultimately used to predict and explain behavior in the stock market, housing market, and employment market. The influence has also extended beyond applied economics to the fundamental core of theoretical economics. A recent review in a prominent economics journal, for example, advised economists to broaden their theories beyond the assumptions associated with "Chicago man" (the rational actor associated with the free-market economic theories developed at the University of Chicago) to incorporate the constraints associated with "K-T man" (McFadden, 1999).

A second boost to the heuristics and biases program is one we have already mentioned, the set of theories and metaphors associated with the "cognitive revolution" that dominated psychology when Kahneman and Tversky advanced their initial set of heuristics. The set of analogies associated with conceptualizing the mind as a computer is congenial to the idea of subroutines devoted to assessments of similarity, availability, and adjustment from some handy starting point. The fit is even tighter, of course, if one conceptualizes the mind (as was quite common in the 1970s) as a computer with limited processing capacity. Such a view makes the idea of effort-saving subroutines that sometimes provide reasonable but imperfect solutions seem particularly appealing and compelling. Sloman (1996) discusses the even closer fit of the heuristics and biases approach with the more modern conception of the mind as a connectionist computer, characterized by massively parallel processing and coherence-based computation (Sloman, Chapter 22, focuses on psychological evidence rather than computational principles).

The heuristics and biases message also fit well with – and was reinforced by – the pragmatic agenda of much of the field of social psychology. Social psychologists have had an enduring interest in social problems and their alleviation. Research on such topics as persuasion, conformity, and cognitive consistency has been fueled by a concern with the dark side of each – sinister propaganda, mindless conformity, and the biases to which rationalization gives rise. But the social evil with the greatest fascination for social psychologists has always been

the combination of stereotyping, prejudice, and discrimination, topics to which the heuristics and biases agenda was seen as highly relevant. Anyone interested in the false beliefs that characterize many stereotypes is likely to be receptive to new ideas about sources of error and bias in everyday judgment.

The field of social psychology was thus receptive to Kahneman and Tversky's ideas from the very beginning and the field's enthusiasm provided another boost to their approach. This is exemplified most powerfully by Nisbett and Ross's (1980) influential treatment of the difficulties people confront in trying to negotiate the complexities of everyday social life, and the nonoptimal strategies they often pursue in the attempt to do so. Their work, which has been called the "errors and biases" perspective in social psychology, was different from Kahneman and Tversky's in an important respect. Nisbett and Ross and their school have been primarily concerned with the causes and consequences of nonoptimal reasoning in social life. Thus, the "fundamental attribution error" (Ross, 1977), the self-serving bias in attribution (Miller & Ross, 1975), and the confirmation bias in social interaction (Snyder & Swann, 1978; Word, Zanna, & Cooper, 1974) have been studied because of their implications for such problems as intergroup conflict and discrimination. In this case, the errors and biases are central; they are not studied first and foremost as a cue to the underlying processes of judgment. (This tradition is developed further by Pronin, Puccio, and Ross in Chapter 36.)

The heuristics and biases message was not only lucky with its supporters, it was also well packaged. Demonstration studies were designed as much like cocktail party anecdotes as traditional cognitive psychology studies, making them magnets for academic lecturers and textbook writers alike. Scenarios involving feminist bank tellers and African countries in the United Nations made the lessons of the heuristics and biases tradition memorable for students at all levels. It is difficult to overestimate the impact of style in the program's success – although the message would not have spread without substance as well. A medium of communication that included stories and personality sketches was well-suited to the message that people think more naturally in terms of narratives and stereotypes than set-theoretic concepts.

CRITIQUES AND CONTROVERSIES

The profile of any intellectual idea is also raised by the controversies it inspires, and the heuristics and biases tradition has inspired many. People, particularly academics, do not accept new ideas and approaches easily, nor should they. As Galbraith noted, "Faced with the choice between changing one's mind and proving that there is no need to do so, almost everyone gets busy on the proof." So it has been with the reaction to the heuristics and biases program – many minds have been busy defending the rationality of everyday judgment and proving that the core ideas of the heuristics and biases program are misguided. Here are the central ideas of some of those proofs.

The "We Cannot Be That Dumb" Critique. The most common critique of the research on heuristics and biases is that it offers an overly pessimistic assessment of the average person's ability to make sound and effective judgments. People by and large manage their lives satisfactorily, something they would be unlikely to accomplish, the argument goes, if their judgments were so prone to bias. Indeed, working collectively, humans have split the atom, recombined DNA, and traveled to the moon. Critics see the heuristics and biases program as denigrating "human decision makers as systematically flawed bumbler" (Ortmann & Hertwig, 2000) because "actual human reasoning has been described as 'biased,' 'fallacious,' or 'indefensible'" (Gigerenzer, 1991a, p. 259). As an outraged team of psychologists queried, "Are heuristics-and-biases experiments cases of cognitive misers' underachieving, or of their receiving a Bayesian hazing by statistical sophisticates?" (Barone, Maddux, & Snyder 1997, p. 143).

This critique owes much of its pervasiveness and appeal to the fanfare that the negative message of the heuristics and biases program has generated at the expense of its positive counterpart. There is, of course, some inevitability to this: Negative information typically dominates the positive. Just as there is a "bad news bias" in media reporting ("if it bleeds, it leads"), it is hardly surprising that the negative message of the heuristics and biases program would capture more attention, inspire more like-minded research, and serve as the focal point of disagreement and controversy. Nonetheless, the common belief that examples of human error are disproportionately cited in the scholarly literature turns out to be an oversimplification; the prominence of such demonstrations is accounted for by the prominence of the journals in which they are found (Robins & Craik, 1993).

There is, however, one version of this critique to which researchers in the heuristics and biases tradition must plead "no contest" or even "guilty." This is the criticism that studies in this tradition have paid scant attention to assessing the overall ecological validity of heuristic processes. Ecological validity (Brunswik, 1955) corresponds to the correlation of the actual outcome in the world with the cue available to the perceiver across a universe of situations. Thus, assessing the ecological validity of the representativeness heuristic would involve identifying a universe of relevant objects (e.g., every scholar in the engineering and law faculties at a given university), and then correlating the outcome value for each object (e.g., membership in either faculty) with the value of the cue variable for each object (e.g., relative similarity to the prototype of each faculty). This correlation, then, would provide a measure for the given universe of how well the representativeness cue performed. This Herculean task has not attracted researchers in the heuristics and biases tradition; the focus has been on identifying the cues that people use, not on evaluating the overall value of those cues. Nevertheless, researchers in this tradition clearly share a set of assumptions: the ecological validities are probably high, the heuristics are generally useful, but common and profoundly important exceptions are to be found. (Note how this summary could be applied to the "fast and frugal"

decision heuristics discussed in Chapter 31 by Gigerenzer et al., despite the apparent opposition between the ecological rationality movement and the heuristics and biases perspective.)

Thus, although there is doubtless some scorekeeping with respect to instances of sound and unsound judgment, it is not of the "box score" sort in which a tally is kept of the number of instances in which people exhibit biased and unbiased judgments. Such a tally is beside the point. A meaningful overall characterization of the quality of human judgment is neither possible nor sought after. To the extent that any such characterization is possible, it would be hard to resist the conclusion that the glass is both half full and half empty. People make a great many judgments with surpassing skill and accuracy, but evidence of dubious belief, questionable judgment, and faulty reasoning is not hard to find (Dawes, 1988; Gilovich, 1991; Schick & Vaughn, 1999; Stanovich, 1986).

Note that there is more than a little irony in the strong form of this critique, one that advances the rather Panglossian notion that people's judgments are hardly ever biased (see Stanovich & West, 2000; Chapter 24, for a consideration of this view). The same scientists who advance this claim use a variety of methodological safeguards such as double-blind experimental designs to make sure their own observations are not contaminated by bias. Are the observations of scientists so much more prone to bias than the individuals they study?

Advocates of the "people are not that dumb" critique have found their voice among evolutionary psychologists for whom it is axiomatic that people perform all tasks critical for survival and reproduction extremely well. According to this school, ancestors who could not reliably make judgments important to survival did not survive long, and therefore the biological basis of their judgmental tendencies have been driven from the gene pool. There is, of course, considerable merit to this perspective. Only a creationist would maintain that our mental faculties were sculpted by something other than evolution. It is also the case that some judgments strike us as hard and others easy, and it is a good bet that the ones that strike us as hard were not subject to the same intense evolutionary pressures as those that strike us as easy. The problems our ancestors absolutely had to solve offer little challenge to us now because of the mental mechanisms we inherited to solve them.

But this logic hardly implies that there is no room for systematic error in judgment. Evolutionary pressures acting on the bulk of human judgments are neither sufficiently direct nor intense to sculpt the kind of mental machinery that would guarantee error-free or bias-free judgment. As Simon pointed out long ago (1956, 1957), evolutionary pressures only lead to local ("better than"), not global ("best possible") optimization. Evolutionary pressures lead to adaptations that are as good or better than a rival's; they do not lead to adaptations that are optimal. If they did, warblers would not rear cuckoo chicks (which they do even though the cuckoo chick is much bigger than the adult warbler), lions would not stalk upwind of their prey (which they do despite greater success

stalking downwind), and people would not probability match in so many different domains (which they do despite paying a high price in foregone gains).

It is ironic that the heuristics and biases approach would be criticized as inconsistent with the dictates of evolution because it *is* an evolutionary account (see Chapter 2 in particular). It is an evolutionary account that recognizes the constraints imposed by an organism's evolutionary history, constraints that yield noteworthy imperfections in function. As Gould (1997, p. 52) argued, "even the strictest operation of pure Darwinism builds organisms full of non-adaptive parts and behaviors. . . . All organisms evolve as complex and interconnected wholes, not as loose alliances of separate parts, each independently optimized by natural selection." The heuristics and biases approach takes the notion of such historical constraints seriously and examines the imperfections that both reveal that history and illuminate current mechanisms.

Kahneman and Tversky's frequent analogies between perceptual and cognitive processes highlight this historical emphasis and reflect an important recognition that cognition evolved after (and out of) perception. Organisms must perceive and act before – or more pressingly than – they need to think, and this doubtless has implications for the quality with which these functions are carried out. Compare, for example, the quality of your motor memory with the quality of your semantic memory. Compare how easy it is to remember how to ski decades after last doing so with how easy it is to recall the trigonometric functions you learned in high school, the foreign language you learned as a teenager, or even all of your childhood romances.

It is clear, then, that there is no deep-rooted conflict between an evolutionary perspective on human cognition and the heuristics and biases approach (Samuels, Stich, & Bishop, in press). Both are concerned with understanding the psychological mechanisms people employ to solve a variety of important real-life problems. Both acknowledge that many cognitive problems essential to survival are typically solved with efficiency and precision. And both can accept the existence of pockets of (particularly informative) bias and error in human judgment.

Indeed, even one of the more popular metaphors of the evolutionary approach to reasoning – that of the mind as a Swiss Army knife – is entirely consistent with the heuristics and biases approach. Although psychologists and neuroscientists have no handle on just how modular the mind might be (Fodor, 2000), it is certainly not unreasonable to suppose that many higher-order cognitive functions are indeed performed by discrete modules. There might be, for example, a module that computes similarity between entities, another that performs basic counting and frequency functions, another that handles causal relations, and so on. Such a mind – one that used different "tools" to perform its various tasks – would produce a pattern of judgments that corresponds perfectly to that documented in the heuristics and biases tradition. At some times and in some contexts, tasks are performed by just the right module and sound judgments are made. At other times and in other contexts, however, specific

tasks are coopted by the wrong module and systematically biased judgments are the result. On still other occasions, of course, the mind might not have the right module to handle the problem (no Swiss Army knife does everything an outdoorsman needs done) and so the task is assigned to a "next best" module, and imperfections in judgment should once again be the result. A modular mind should also produce a pattern of judgments whereby a problem described or structured in one way yields one type of response, whereas the same problem described or structured another way yields a vastly different response – exactly the pattern of results reported countless times in the heuristics and biases literature.

The "It's All Parlor Games" Critique. Another common critique of the heuristics and biases tradition has been to dismiss the reported findings as mere laboratory curiosities – as demonstrations that people cannot readily solve tricky "word problems." The implication is that judgment outside the laboratory is likely to look far superior to that exhibited within.

This critique overlooks that it was the existence of biased judgments in the real world that motivated the heuristics and biases research program. Recall that an important impetus for this research was the work by Paul Meehl on the problems inherent in expert clinical judgment. Recall also that it was the observation of faulty reasoning among students trying to learn statistics (e.g., the gambler's fallacy, the regression fallacy, insensitivity to sample size) that gave the research its initial shape. This critique also flies in the face of the influence that the heuristics and biases research program has had across a wide range of applied disciplines, something it could not do if it dealt only with contrived, artificial problems. As we have noted, the heuristics and biases program has influenced scholarship and curricula in political science, medicine, law, and management.

One particularly persistent form of this critique is the claim that the biases revealed in this research are merely the product of fiendish (or clueless) experimenters who ask misleading questions. Participants are not responding incorrectly, in other words; they are giving the right answer to a different question than the one the experimenter believes he or she is asking.

There is doubtless some merit to this claim, at least as applied to some individual experiments that purport to demonstrate a given bias or shortcoming of human judgment. There is a complex psychology – a subtle set of tacit assumptions and implicit demands – that accompanies participation in a psychology experiment. Even investigators attuned to this psychology can sometimes fail to anticipate correctly how a stimulus is likely to be construed or a question interpreted by a given participant. It is no small task for experimenters to "get it right," which is why psychological research requires so much pretesting.

But just as it is clear that some individual experiments are open to this critique, it is equally clear that the main biases uncovered in this research tradition (e.g., availability biases in likelihood estimates, insensitivity to sample size and prior probability, the conjunction fallacy, anchoring, packing and unpacking

effects) are not. These have all been demonstrated in countless contexts and with varied paradigms and dependent measures, and with domain experts as well as student volunteers. Although an isolated demonstration of some of these biases may be open to this critique, the overall body of evidence in support of them is not. For example, in one of the original demonstrations of the conjunction fallacy using the famous "Linda problem," it is entirely possible that participants interpreted the option "is a bank teller" to mean "is a bank teller who is *not* active in the feminist movement" given that one of the other options was "is a bank teller who is active in the feminist movement." Participants who construed the former this way can hardly be faulted for ranking it lower in likelihood than the latter. But this alternative interpretation simply cannot handle the observation that when participants in a between-subjects design rated (rather than ranked) the likelihood of only one of these options, those evaluating Linda the feminist bank employee offered higher likelihood ratings than those evaluating Linda the bank employee. In such a between-subjects format, of course, there is no artifactual basis for participants to conclude that "bank teller" is to be interpreted as "nonfeminist bank teller." To the extent that any participants did so, it was because their mental model of bank teller crowded out the possibility of feminism, not because they could infer that the experimenter intended "bank teller" to mean "nonfeminist bank teller." (Chapters 1 and 2 contain more detailed discussions of alternative interpretations of the conjunction fallacy.)

This example raises the important point that certain misconstruals on the part of participants are not artifacts, they *are* the phenomena of interest. There is a long tradition of research in social psychology illustrating that people actively construe the meaning of a given task or stimulus (Griffin & Ross, 1991) and that their own chronically accessible categories, habits, and experiences powerfully influence their construals (Higgins, King, & Mavin, 1982; Higgins, Rholes, & Jones, 1977; Srull & Wyer, 1979, 1980). Consider Kahneman and Tversky's well-known engineer/lawyer problem. When asked whether a given description is likely to belong to an engineer or lawyer, one cannot fail to compute the similarity between the description and each professional stereotype. Both the immediacy and relevance of this assessment of similarity, then, make it highly likely that one's very definition of what the task is about will be hijacked. A question about likelihood is construed as a question about "fit." And there is nothing artifactual about the switch (on this point, see Kahneman & Frederick, Chapter 2; Bar-Hillel and Neter, Chapter 3, this volume). It happens both inside and outside the laboratory, and, just as one would expect if people were active interpreters of the tasks that confront them, various changes in the presentation of stimuli or the description of the task influence what participants interpret their task to be (e.g., Koehler, 1996; Macchi, 1995).

The "It's Not an Error" Critique. Another common accusation against the heuristics and biases tradition is that researchers hold experimental participants to an inappropriately high or even misguided standard of rationality. Jonathan

Cohen, for example, contends that rationality does not exist apart from human intuition (Cohen, 1981) – in fact, standards of rationality are the "distillation of our best intuitions about rationality" (Papineau, 2000, p. 173). Thus, how can the source of our standards of rationality prove to be irrational? Perhaps people – especially people participating in unfamiliar or otherwise misleading experimental games – make performance mistakes that mask their underlying rational competence, but by definition, human intuition must be rational.

This critique usefully points out two aspects of the "rationality problem." First, it has a distinct "Alice in Wonderland" flavor: People can and do define rationality in many contradictory ways (see Evans and Over, 1996, for an attempt to deal with the definitional problem). Second, it brings to the fore the crucial role of axioms in justifying a normative theory. Probability theory is built from a few fundamental rules or axioms that reflect people's considered opinions in the abstract. These axioms include the claim that the probability of an event *A* and its complement *not-A* sum to 1. This is an example of a coherence axiom that constrains the relation between the probabilities of events. Axioms have force only if people agree with them – it is possible to opt out and agree not to be bound, but most people find such rules to be compelling. In fact, it is the tension between the general agreement with the abstract rules of probability and the violation of those rules in richer contexts that give the heuristics and biases demonstrations their power (a point explored more deeply in Chapters 1, 2, and 22).

A number of prominent statisticians and philosophers have opted out from any version of probability theory that deals with unique or "one-shot" events. Supporters of the "objective frequency" or "relative frequency" approach (e.g., von Mises, 1928) restrict the domain of probability theory to events that can be repeated in an infinite series. According to this perspective, probability is defined as the relative frequency of an event in an infinite series. Other types of events – such as predictions, beliefs, or statements about a single case – cannot be evaluated by the rules of probability theory. There is no denying the appeal of dividing the domain of uncertainty into the "objective" (like the spin of a roulette wheel or the toss of a coin) and "subjective" (like a prediction of tomorrow's weather). However, as Keynes (1921) argued, a strict frequentist view entails that beliefs about unique events such as the coming of war, the end of a recession, and the outcome of a medical operation cannot be evaluated. And even those who take the frequentist stance in their professional lives act like subjectivists in their day-to-day affairs. An honest frequentist must concede that meaningful probabilistic statements can be made about unique events, such as the Yankees being more likely to win the World Series this year than, say, the Kansas City Royals or Montreal Expos, or that either author of this chapter is likely to lose a prize fight with the reigning champion in his weight class.

Such consideration notwithstanding, at some point one is thrown back to the level of axioms: Is one willing, for example, to subscribe to the rule that a set of predictions made with 80% probability should come true 80% of the time?

For those willing to opt in to such a “correspondence axiom,” the demonstrations found within this book should be particularly relevant. However, even those readers who by personal taste or ideological alliance reject the correspondence axiom might be swayed by the “expected value” argument as characterized by the statistician De Finetti. “The formal rules normally used in probability calculations are also valid, as conditions of consistency for subjective probabilities. You must obey them, not because of any logical, empirical or metaphysical meaning of probability, but simply to avoid throwing money away” (De Finetti, 1970, p. 137). Finally, those frequentists who are unwilling to have their own beliefs evaluated by the coherence or correspondence axioms may still be curious to find out how well the classical rational actor model – incorporating the axioms of subjective probability – stands up to empirical test.

The “Frequencies, Good; Probabilities, Bad” Critique. Given the controversy surrounding the normative status of frequencies and subjective probabilities, it is not surprising that those who favor an evolutionary defense of rationality (“ecological rationality”) should throw in their lot with the frequentists. Evolutionary psychologists (e.g., Pinker, 1997) maintain that success in our ancestral environment required only a talent for working with frequencies, not probabilities. This argument, precisely because it cannot be tested empirically, remains a matter of faith and ideology. However, the frequentist argument for evolutionary rationality contains a component that can be tested empirically: The evidence for heuristics and biases, it is claimed, “disappears” when stimuli are presented and questions are asked in terms of frequencies (Gigerenzer, 1991b; 1994).

This was a bold argument when first introduced and it is even bolder to maintain now (e.g., Gigerenzer, 1998; Cosmides & Tooby, 1996) when a score of studies have indicated that it simply does not hold up empirically. In fact, presenting frequencies rather than probabilities sometimes makes judgment distinctly worse (e.g., Griffin & Buehler, 1999; Treadwell & Nelson, 1996), sometimes makes judgments distinctly better (e.g., Tversky & Kahneman, 1983; Koehler, Brenner, & Tversky, 1997) and quite often leaves the quality of judgment largely unchanged (Brenner, Koehler, Liberman, & Tversky, 1996; Griffin & Buehler, 1999). Even more troublesome for the evolution/frequency argument, Kahneman and Tversky’s original explanation of the probability–frequency discrepancy (Kahneman & Tversky, 1982a; Tversky & Kahneman, 1983) provides a unified account of when frequency formats improve judgments and when they do not (e.g., Sloman, Slovak, & Over, 2000).

Critics claim that assessments of single-event probabilities are unnatural, and that only a frequency format is consistent with how the mind works (Cosmides & Tooby, 1996; Gigerenzer, 1991b, 1994; Pinker, 1997). Kahneman and Tversky argued, in contrast, that representing problems in terms of frequencies tends to evoke mental models that facilitate the detection of set inclusion relations and thus improves judgment – and this view has received considerable support from the studies of Sloman and others (e.g., Evans, Handley, Perham, Over, &

Thompson, 2000; Girotto & Gonzalez, 2001; Sloman & Over, in press; Sloman et al., 2000).

Note that even some scholars who are sympathetic to the heuristics and biases tradition have concluded that its lessons are limited to “problems with probabilities.” “While we have had the intellectual resources to pursue truth for at least 100,000 years, and quite possibly a lot longer, the notion of probability has only been around since 1654. . . . It is no accident that most of the ‘irrationality’ experiments trade in probabilities” (Papineau, 2000, p. 182). Thus, it is important to point out that although frequentistic formats – for whatever reason – sometimes induce more effective processes of judgment, it is simply not the case that the biases uncovered by Kahneman, Tversky, and others “disappear” if people are allowed to think in terms of frequencies rather than probabilities. Numerous studies in the heuristics and biases tradition make this clear, any one of which is sufficient to make the point. For example, in one of the earliest of Kahneman and Tversky’s experiments, participants were asked to estimate either the number of possible committees of 2 people from a group of 10, or the number of possible committees of 8 from a group of 10. Here, as in many other studies, the participants were given no information in a probabilistic format nor was a probabilistic response required (or even possible). Nevertheless, even though the actual number of possible committees of 2 and 8 are the same, those estimating the number of 2-person committees gave estimates that were an average two and a half times higher than those estimating the number of 8-person committees (Tversky & Kahneman, 1973). There is clearly more to biased judgment than an inability to handle probabilities.

Recent Perspectives

As this book demonstrates, the heuristics and biases program has weathered its various critiques and remains a vigorous and still developing perspective on human judgment. Part of its vigor stems from parallel developments in psychology that have both influenced and been influenced by the work on heuristics and biases. Most significant in this regard has been the broad interest in effortless or “automatic” mental processes that play an important role in a wide variety of everyday phenomena. Work on the rapid, automatic assessments of the affective system is a good example (Slovic, Finucane, Peters, & MacGregor, Chapter 23; Zajonc, 1980). The idea that a quick, affective, yes/no, approach/avoid reaction precedes extensive cognitive elaboration has certainly been around for a long time and predates the heuristics and biases program. Current thinking about the way such a process operates, however, has been shaped by the heuristics and biases program, a result seen most clearly in the use of such terms as Slovic and colleagues’ “affect heuristic” (Chapter 23) and Pratkanis’s “attitude heuristic” (Pratkanis, 1989). Contemporary research on magical thinking also fits well with the heuristics and biases perspective (see Chapter 11 by Rozin & Nemeroff). This research, like that in the heuristics and biases tradition, highlights the conflict between an initial, reflexive evaluation

and a more considered, rational assessment. The idea that heuristics often operate automatically is also compatible with current research on “automaticity” (Bargh, 1997; Bargh & Chartrand, 1999). This work, like the heuristic and biases program, stresses the fact that much of mental life is not the product of deliberate processing, but of quicker, more reflexive processes that are less available to conscious intervention.

Much of the work on relatively effortless, reflexive mental processes that followed in the wake of the research on the heuristics of judgment has been advanced as part of various “dual process” models of cognition. The advocates of each of these models postulate one set of mental processes that are quick and effortless, and another that are more deliberate and taxing. There are two types of dual process models. One advances the claim that people deliberately use less effortful procedures when the judgment is relatively unimportant and motivation is low. The more effortful procedures are reserved for occasions in which the stakes are high. The “Elaboration Likelihood” (Petty & Cacioppo, 1986) and “heuristic-systematic” models of persuasion (Chaiken, Liberman, & Eagly, 1989), and various perspectives on stereotyping (Bodenhausen, 1990; Fiske & Neuberg, 1990) fit this template. This work fits the “cognitive miser” perspective on heuristic processing.

There is another set of dual-process models that do not conform to the cognitive miser perspective. These models, often referred to as “two systems” models, postulate the existence of two mental systems that operate in parallel. (See Kahneman & Frederick, Chapter 2; Slovic, Chapter 22; Stanovich & West, Chapter 24.) An associationist, parallel-processing system (“System 1”) that renders quick, holistic judgments is always in operation – not just when motivation is low and judgments are made on the cheap. The assessments made by the associationist system are then supplemented – and sometimes overridden – by the output of a more deliberate, serial, and rule-based system (“System 2”). These models fit the cognitive miser perspective less well because they do not postulate two different “routes” of information processing that operate in either-or fashion according to the motivation of the information processor (although they too can account for motivational influences through variation in the effort applied to the rule-based system, e.g., Tetlock, Chapter 32).

As we alluded to earlier, the heuristics and biases program has most often been seen through the cognitive miser lens. People are thought to employ various heuristics to save effort. But the idea of heuristics as “natural assessments” (Tversky & Kahneman, 1983; see Chapter 1) is clearly much more consistent with the two-systems perspective, and we expect this book to consolidate that view of what the heuristics and biases program is really about. The two-systems view also helps to clarify the differences and similarities between the heuristics of the “heuristics and biases” program and those of the “fast and frugal heuristics” program (Gigerenzer et al., Chapter 31; Gigerenzer, Todd, & the ABC Research Group, 1999). The prototypical fast and frugal heuristics studied by the ABC

group are System 2 heuristics: strategies or rules that are deliberately chosen to ease computational burden (Kahneman & Frederick, Chapter 2; Griffin & Kahneman, in press). The “Take the Best” heuristic, for example, simplifies multiattribute choice by using an ordinal comparison on the most important dimension.

Some of these System 2 heuristics, however, rest on more basic System 1 processes that are subject to the kinds of errors associated with the heuristics and biases approach. Consider the “recognition heuristic,” by which people choose the option that is most familiar. Familiarity appears to rely on the computation of fluency – how easily an object is identified or a word is read (Kelley & Jacoby, 1996) – which is an automatically computed natural assessment closely related to availability (Schwarz, Chapter 29). Thus the accuracy of the (System 2) decision rule rests on the validity of a System 1 computation, which is sensitive to a wide variety of environmental manipulations that lead to robust and systematic biases (e.g., Kelley & Jacoby, 1996; Reber, Winkielman, & Schwarz, 1998). It seems that System 1 heuristics, and the biases that arise from them, are difficult to avoid even in the context of deliberate choice (Frederick, Chapter 30; Over, 2000).

What are the natural assessments that are automatically elicited by certain stimuli or certain judgment contexts? In addition to the computations of similarity and availability that were the basis of the original research in this tradition, people typically make quick assessments of their affective reactions (Schwarz, Chapter 29; Slovic et al., Chapter 23), the ease or fluency of their perceptual experience (Kelley & Jacoby, 1996; Reber, Winkielman, & Schwarz, 1998), the causal structure of the pertinent environment (Heider, 1958; Michotte, 1963; Tversky & Kahneman, 1982) and whether a given event is abnormal or surprising (Kahneman & Miller, Chapter 20; Kahneman & Varey, 1990). There may very well be other natural assessments that constitute additional general-purpose heuristics that have yet to be empirically documented. In addition, there are certain assessments that are only “natural” in some contexts and therefore serve as the basis of various special-purpose heuristics. In some contexts, for example, a person is likely to compute how hard she has worked to get something, and then use the output of that computation to infer how much she values it (Aronson & Mills, 1959; Bem, 1972; Gerard & Mathewson, 1966).

The assessments underlying the six general purpose heuristics identified (affect, availability, causality, fluency, similarity, and surprise) are different from the two discussed most actively in 1982 when the predecessor of this book was published. The notion of special-purpose heuristics, absent from the heuristics and biases landscape at that time, is even more novel (Kahneman & Frederick, Chapter 2; Frederick, Chapter 30). We view these new developments as a sign of the vigor and, to use the other sense of the term, heuristic value of the heuristics and biases perspective. Part I of this book presents papers that develop classic themes in this perspective, revealing how the original roots of the research

continue to deepen. Part II presents papers that provide new theoretical approaches to heuristics and biases, revealing how the roots continue to spread. Finally, Part III presents papers that describe judgment in everyday life and by domain experts, showing how the heuristics and biases research has branched out into applied settings. At every level, research inspired by this program is flourishing, and we look forward to seeing what will grow out of the work summarized in this book.

A REPRESENTATIVENESS AND AVAILABILITY

PART ONE

THEORETICAL AND EMPIRICAL EXTENSIONS

Uncertainty is an important part of life. Choices must be based on incomplete information, the guilt of a defective decision, the uncertainty of a defective outcome. In such situations, people normally do not have access to the rules that govern such events, intuitive judgments are often used to assess uncertainty.

The question of how to assess the probability of certain events has always been a central theme in the history of thought (Hogarth, 1981; Kahneman, 1981). One of the central themes of this research has been the extent to which people's judgments of the probability of events conform to the rules of statistics and probability theory. People do not use the probability rules of statistics and probability theory in their judgments. Perception might compare the size of an event to the size of the sample space. Unlike the correct size of an event, the size of the sample space is not easily defined. Because people hold different beliefs about the size of the sample space, no single sample space probability cannot always be defined. The domain of random sampling is the domain of uncertain events. People's judgments of uncertain events are often among them. For example, if A is a subset of B, then the probability of A must be less probable than the probability of B.

The laws of probability do not apply to events whose measure is defined on a random process. The measure is defined on a random process of possibilities, such as the chance of a coin landing heads or tails or the chance of a die. The probability of an event is the number of favorable outcomes divided by the number of possible outcomes. Probability theory applies to repetitive chance processes, but the laws of probability do not apply to unique events for which probability is not defined. The probability of a "favorable" outcome is the number of favorable outcomes divided by the number of possible outcomes.

This is an extended version of a paper that originally appeared in *Psychological Review*, 1982, 89, 1-14. Copyright © 1984 by the American Psychological Association. Adapted with permission.