Judgment Under Uncertainty: Heuristics and Biases.

Daniel Kahneman, Paul Slovic, and Amos Tversky (eds.). New York: Cambridge University Press, 1982. xiii + 555 pp. $44.50 ($14.95 paperback).

This remarkable volume draws together the most important results of a decade's research into how people make judgments of probability and how, more generally, they make judgments under uncertainty. It includes the work of 39 authors, most of them psychologists. It emphasizes, quite appropriately, the work of its editors. Of the 35 papers, 12 are by Kahneman and Tversky and four more are by Slovic and his associates at Decision Research in Oregon.

The volume is organized into ten parts. The following is a list of some of the themes covered.

Part I. *Introduction.* Here we find Tversky and Kahneman (1974), which reviews three heuristics people appear to use in making judgments of probability: representativeness, availability, and anchoring and adjustment.

Part II. *Representativeness.* When trying to assess the probability that A is a B, people often resort to assessing how similar A is to the typical B. Asked how probable it is that a given person is a librarian, we may think about how similar he is to our stereotype of a librarian. This can lead to biases: We lose sight of how few people are librarians. People are also misled by the representativeness heuristic when thinking about random sampling—more representative samples seem more likely.

Part III. *Causality and Attribution.* Our everyday thinking is heavily causal, and we tend to weight causal arguments (scenarios, attributions of motives) more heavily than more abstract evidence (approximately known frequencies or "base rates"). Several of the papers in this section are by social psychologists.

Part IV. *Availability.* When we try to assess the frequency of an event by sampling from our memory or imagination, our judgment may be affected by the differential availability of examples. Some things are more easily retrieved from memory than others; some scenarios are more easily constructed than others. Again we see examples from social psychology. A person may overestimate his relative contribution to a joint undertaking because examples of his own efforts come to mind most readily.

Part V. *Covariation and Control.* People are not able to discern weak ($r < .8$) correlations in empirical data. They are prone to see nonexistent correlations that are suggested by verbal associations.

Part VI. *Overconfidence.* Here we find two classics: Oskamp's (1965) demonstration that clinical psychologists are overconfident in their predictions of patients' behavior, and the frequently cited but not previously published report by Alpert and Raiffa (1969) on people's tendency to construct excessively tight probability distributions. We also find a review of the literature on calibration and an essay on the mechanisms of hindsight.

Part VII. *Multistage Evaluation.* People tend to underestimate the result of multiplying many moderately sized factors. Consequently, they tend to overestimate the probability of a sequence of events that is expressed as the product of many fractions. In contexts that are not overtly probabilistic they also tend to ignore uncertainties in the less salient steps of a sequence. These tendencies can lead to systematic biases in the estimation of likelihood ratios.

Part VIII. *Corrective Procedures.* How can people improve their probability judgments? Simple arithmetic can help. Training in statistical theory can help; people can learn to take sample size into account and to shrink subjective impressions towards the mean. But many of the techniques suggested by statistical theory cannot, because of lack of information, be implemented in everyday reasoning.

Part IX. *Risk Perception.* The perception of risk shows systematic biases. For example, people tend to overestimate the risk of death from catastrophic accidents and underestimate the risk from common disease. The biases that affect judgments of probability in general affect both lay and expert assessments of technological risk.

Part X. *Postscript.* Kahneman and Tversky address methodological and conceptual issues that have been raised by critics and sketch a classification for types of uncertainty.

Heuristics and Biases

The word "heuristic" is used very broadly in this book. When the authors call a given method a heuristic, they do not mean that there is a more reliable method. And when they say people usually use heuristics, they do not mean that people usually make probability judgments hastily and carelessly.

The word "bias" is also used very broadly here. The heuristic of representativeness, for example, is called biased because it can be improved by shrinking its predictions towards the mean. Statisticians are accustomed to a narrower usage. They call the sample median an unbiased estimator of the center of a normal distribution, even though it can be improved by being shrunk toward the sample mean.

The authors might be better understood if they avoided the terms "heuristic" and "bias." Instead of saying that people's heuristics are sometimes biased, they might simply say that people's methods can sometimes be improved. This milder language might forestall some of the strictures of philosophers who feel that psychologists have been too quick to call people irrational. (See, for example, Cohen 1981.)

Why Should Statisticians Care?

There are a number of reasons why this volume should be of interest to statisticians.

First, as teachers of statistics we need to understand how people approach problems of probability when they do not understand or fail to use the theory we are trying to teach. A student may find more meaning in the rules of probability after he has confronted examples where his untutored judgments violate these rules. And he may find it easier to recognize when simple statistical techniques are applicable if he has compared these techniques to the methods he would use if he did not have them.

Second, statisticians need to learn about the pitfalls in people's intuitive reasoning about probability because statisticians are people too. Consider Oskamp's study of overconfidence among clinical psychologists. When working with fragmentary information, psychologists make fairly accurate estimates of the low reliability of their predictions about a person's behavior. As information accumulates the accuracy of their predictions scarcely improves, but their confidence in their predictions increases dramatically. Few would deny that statisticians often fall into the same trap. As we elaborate a model and gather more data to estimate its parameters, we can easily lose sight of the limited reliability of the model for the actual questions we are answering.

Third, and most important, an understanding of what people do (and of what they can do) when asked to make subjective probability judgments is essential to wider implementation of theories of subjective probability.

The Bayesian theory provides "normative" constraints on subjective probability judgments, and Bayesian statisticians stress that the failure of people to obey these constraints does not make the constraints any less normative. Other theories of subjective probability have often been given a similar "normative" interpretation. But none of these theories has explained how people should go about making judgments. The work presented in this volume presents a challenge to these theories because it demonstrates that the heuristics people have available for making intuitive probability judgments cannot be expected to automatically produce judgments that come anywhere close to satisfying the normative constraints. Clearly theories of subjective probability need to become less normative and more "constructive"—they need to spell out how an individual is to construct subjective probability distributions, "belief functions" (Shafer 1982), or "possibility distributions" (Zadeh 1978). Presumably the first steps in this construction will involve the use and adaptation of the heuristics people are known to be able to use.

Related Work and Literature

The editors of this volume have also made very important contributions to the study of how people construct values and make choices and decisions (see especially Kahneman and Tversky 1979; Fischoff, Slovic, and Lichtenstein 1980; and Tversky and Kahneman 1981). They do not include any of this research in this book. "The topic of decision making," they write in their preface, "is important enough to be the subject of a separate volume." This separate volume will be eagerly awaited.

The most important related work to appear recently, Tversky and Kahneman (1983), concerns the "conjunction fallacy."

Conclusion

This is an important and well-written book. The paperback edition, at least, is well worth the price. It belongs in the library of everyone who teaches or practices probability judgment.

GLENN R. SHAFER
University of Kansas

REFERENCES

ALPERT, M., and RAIFFA, H. (1969), "A Progress Report on the Training of Probability Assessors," Unpublished.

COHEN, L.J. (1981), "Can Human Irrationality be Experimentally Demonstrated? (With Discussion)," *The Behavioral and Brain Sciences*, 4, 317–370.

FISCHOFF, B., SLOVIC, P., and LICHTENSTEIN, S. (1980), "Knowing What You Want: Measuring Labile Values," in *Cognitive Processes in Choice and Decision Behavior*, ed. T.S. Wallsten, Hillsdale, N.J.: Erlbaum, 117–141.

KAHNEMAN, D., and TVERSKY, A. (1979), "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 263–291.

OSKAMP, S. (1965), "Overconfidence in Case-Study Judgments," *Journal of Consulting Psychology*, 29, 261–265.

SHAFER, G. (1982), "Belief Functions and Parametric Models," *Journal of the Royal Statistical Society, Series B*, 44, 322–352.

TVERSKY, A., and KAHNEMAN, D. (1974), "Judgment Under Uncertainty: Heuristics and Biases," *Science*, 185, 1124–1131.

——— (1981), "The Framing of Decisions and the Psychology of Choice," *Science*, 211, 453–458.

——— (1983), "Extensional vs. Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment," *Psychological Review*, 90, 293–315.

ZADEH, L. (1978), "Fuzzy Sets as a Basis for a Theory of Possibility," *Fuzzy Sets and Systems*, 1, 3–28.

Expected Utility Hypotheses and the Allais Paradox.

Maurice Allais and Ole Hagen (eds.). Dordrecht, Holland: D. Reidel Publishing, 1979. vii + 714 pp. $86.85.

The early 1950's was a time of great excitement for economists, psychologists, and statisticians interested in the idea of numerical utility. Von Neumann and Morgenstern had only recently, in 1944, published their axiomatization of expected utility. L.J. Savage was at work on their suggestion that utility and subjective probability be jointly axiomatized at a similar level of rigor. The expected utility model seemed to offer exciting possibilities as a descriptive model in psychology, as a foundation for microeconomics, and as a basis for the theory of statistical inference.

For economists, the most important fact about von Neumann and Morgenstern's work was that it revived the idea of cardinal as opposed to mere ordinal utility. The 19th-century marginal utility theorists had talked about utility as if it could be measured on a cardinal scale—that is, as if utility differences could be compared. And the 19th-century German psychologist Fechner had claimed to find experimental bases for defining cardinal utility. But by the early 20th century, cardinal utility had been generally abandoned. Psychologists had become skeptical of Fechner's claims, and economists, under the prodding of Pareto, had recognized that the utilities of the marginal theory had only ordinal meaning. Von Neumann and Morgenstern's utilities, since they were derived in a probabilistic context and scaled by probabilities, did have a cardinal meaning.

In 1952, the French economist Maurice Allais mounted a direct challenge to the new cardinal utility. The axioms of the new theories, Allais declared, were wrong. They were wrong descriptively; people did not obey them. And they were wrong normatively; it was rational to reject them. Allais did believe in cardinal utility; he believed, like Fechner, that it could be defined operationally by direct comparison of differences in utility or scaled by taking minimum perceptible changes as unit changes. But since he rejected von Neumann and Morgenstern's axioms, he saw no reason why a person should rank probability distributions of utilities by expected utility alone. A person might want to consider the shapes of the distributions, perhaps by considering their variances and higher moments. Using various assumptions about how people might rank distributions of utilities, Allais developed his own theory of the economic phenomena associated with risk.

The most effective part of Allais's attack was his empirical demonstration that the tendency to prefer sure gains leads people to express preferences between gambles that are incompatible with expected utility. The idea, roughly, is that people will prefer a 100% chance of winning a large prize to a 99% chance of winning a much larger prize, but will not agree to nearly so great a reduction of the prize in exchange for a single percentage point of probability when that percentage point does not bring them to certainty. This phenomenon has come to be known as "Allais's paradox." For details, see the volume under review, Savage (1954, pp. 101–103), or Raiffa (1968, pp. 80–86).

Though Allais's paradox drew much attention, it did not have the effect that Allais had hoped. During a lunchtime conversation at the 1952 conference where he first presented his ideas, Allais showed his paradox to Savage, and Savage was startled to find himself expressing preferences violating his own theory. But after reflection Savage decided to change his preferences instead of his theory. Expected utility remained, Savage felt, normative. Other Bayesian statisticians have reacted similarly. Economists, who have a greater stake in the descriptive validity of expected utility, have also minimized the importance of Allais's examples, usually by arguing that they involve unrealistically large prizes.

Economists almost completely ignored Allais's own theory. (Some economists, such as Tobin (1958), have considered variances in ranking probability distributions, but the emphasis here has been on ranking distributions of money values, with hope of approximate agreement with ranking by expected utility.) For most economists, Allais's founding of cardinal utility on 19th-century psychology was unacceptable.

The volume under review represents an attempt to reopen the question. It reflects a feeling on the part of Allais and his coeditor, Ole Hagen, that the neglect of Allais's work is due in part to its comparative inaccessibility to those who do not read French. This is remedied here by the publication of 380 pages by Allais: a 120-page translation of his 1953 French monograph (only a 40-page version was published in English at the time) and, further, a 260-page review of the issues. Hagen contributes 30 pages on a theory similar to Allais's. About 40 pages are taken up by critics of Allais, and the remaining 200 pages are devoted to related issues. Altogether 17 authors are represented.

On the whole, the volume is difficult and not as rewarding as the reader might hope. Allais's contributions are well written but far too lengthy, and the quality of the other contributions is uneven. But there are several contributions that do deserve to be flagged:

- Oskar Morgenstern's brief defense of the von Neumann-Morgenstern theory is of historical interest if only because it shows how little objections to the theory impressed him. (He died in 1977.) He defends the descriptive value of the theory, conceding only that it, like Newton's theory, is only approximate and applies only to a limited domain. And he defends its normative status, asserting that people will correct their deviations from it just as they correct mistakes in long division.
- Cyert and DeGroot, in "Adaptive Utility," explore the idea that a person might be unsure about his utilities and use Bayesian methods to learn about them. Since Allais's paradox is sometimes presented as an inconsistency between choices at different times (with both prizes uncertain I will gamble on the larger, but I will choose differently once a favorable event assures me certain success if I opt for the smaller), this might offer a way to accommodate the paradox within expected utility theory. Ironically, though, it is easier to make sense of the idea of unknown utility if we define utility hedonistically, as Allais does, than if we derive it from an extensive set of preferences, as the expected utility theory does.
- Peter Fishburn's elegant contribution is one of his many recent studies of the consequences of relaxing various of the axioms of expected utility theory.
- MacCrimmon and Larson report empirical studies showing deviations from expected utility that go well beyond Allais's paradox. Some of their examples are similar to examples reported by Kahneman and Tversky (1979).

How should we evaluate today Allais's criticism of expected utility theory? To evaluate it adequately we must, I think, consider the full force of the results of recent investigations of the sensitivity of people's choices to context and presentation. Some of these results are reported in Wallsten (1980), Tversky and Kahneman (1981), and Hogarth (1982).

It turns out that choices can be presented in ways that will induce people to violate not only the axioms to which Allais objects but also the axioms to which Allais subscribes, such as the principle of stochastic dominance. In fact, seemingly innocuous changes in wording that do not even objectively change the options in a problem can systematically reverse the choices people make. Tversky and Kahneman (1981, p. 453) give an illuminating example. After describing a public health problem that is expected to cause 600 deaths if unchecked, people are asked to choose between different preventive programs on the basis of their probable consequences. People's choices, it turns out, depend on whether the consequences are described in terms of total deaths or in terms of the number of the 600 saved.

The commonsensical conclusion is that people have neither the extensive preferences required by von Neumann and Morgenstern's theory nor the psychological utilities required by Allais's theory. When offered a choice between gambles they do not query themselves about