

probabilities in the records of the Dutch foreign policy decisions. When Gallhofer and Saris coded the verbal expressions of probability and utility in these accounts, they found an expression of the intensity of both the probability and the utility in only 3 of 235 decisions. Their set of seven decision rules was sufficient to match all but 8 of the 235 decisions, but the SEU strategy was among the least popular of the seven.

Every decision making research group should have and use a copy of the book, and it will be worthwhile in graduate students' reading. The publisher should be faulted for the excessive number of typographical errors in some of the later chapters.

ROBERT M. HAMM

*Department of Psychology,
University of Colorado, Boulder, Colorado, USA.*

THINKING AND DECIDING. Jonathan Baron. Cambridge University Press, 1988.

Jonathan Baron's *Thinking and Deciding* is all about good thinking and poor thinking. The author suggests that all goal-directed thinking and decision making can be described in terms of a *search-inference framework*. That is, thinking can be described as inferences made from possibilities, evidence, and goals that are discovered through searching. This framework is both descriptive and prescriptive. It defines poor thinking as ignoring possibilities, evidence, and goals that should be considered and as making inferences biased toward preserving favored ideas and beliefs. What is needed to overcome poor thinking is active open-mindedness. Clearly this book is not simply a survey but reflects a point of view.

The author defines thinking as a method of choosing among potential possibilities where possibilities might consist of actions, beliefs or personal goals. These choices are based on a search for relevant information and the inferences which are made from the information obtained. Information is broadly defined to include goals, choices, and evidence. Choices are evaluated on the basis of evidence in light of the goals. For example, an undergraduate student having the goal of achieving good grades might rate a course more favorably upon hearing that the average grade in the course is a B. Whether or not taking this course is a good idea might depend upon how superficial the search process was. With further search the student might discover that he or she also has the goal of learning something and obtain evidence that the course will fail to satisfy this goal, or the student might uncover an alternative course where the average grade is an A. Furthermore, the student might determine that although the average grade is a B, the only students who take the course have a natural talent for it, a talent that this hypothetical student might lack. Thinking may also be ineffective if initial ideas bias the search for information or if the search process is too extensive in relation to the importance of a question. Buying a hamburger should not involve a more thorough search than buying a house.

The above analysis of thinking focuses on inadequate

search as the cause of poor thinking. Dr Baron does point out, however, that without *knowledge*, thinking is an 'empty shell.' Because thinking involves search, there must be some contents for the search to operate on.

The author provides a notion of rationality in keeping with his emphasis on goals as organizing search. The best kind of thinking, rational thinking, is 'whatever helps fulfill our personal goals.' Thus violations of the laws of logic and self-deception could, in principle, help fulfill goals and thereby satisfy this definition of rationality.

The book is organized into four main sections. The first section covers general topics concerning thinking and includes chapters on problem solving, simple and complex learning, intelligence, creativity, formal logic, and everyday reasoning. The second section focuses on probability and belief and provides chapters on normative theories of probability, descriptive theories of probability judgement, hypothesis testing, correlation and contingency, and biases and beliefs. The third section covers normative and descriptive theories of utility and choice, quantitative judgement (lens model), impression formation, adding and averaging models), moral thinking, social dilemmas and decisions about the future. Notably, the author's prescriptive point of view extends to moral thinking where he argues against relativism and in favor of a utilitarianism that aims to maximize expected utility (good) for everyone. The book concludes with a chapter on the teaching of thinking and decision making.

The book is intended for advanced undergraduates or beginning graduate students taking a course on decision making and thinking. More broadly, the book is for 'anyone who is disturbed by poor thinking.'

How well does the book succeed in realizing the author's goals? Very well, in my opinion. Each of the chapters is densely packed and never boring. Although the coverage is broad, depth has not been sacrificed. I liked too many of the chapters to single out favorites but here is a small sample of my positive reactions. The chapter on intelligence provides a nice debunking of IQ as *the* true measure of intelligence. The chapter on formal logic presents a clear and thoughtful review in less than 20 pages of text. I did not anticipate that there was much to be said about creativity from the

perspective of the book's goals but was very pleasantly surprised by the author's treatment. Insightful examples and comments truly are too numerous to mention and they range from bringing out the virtues of incorrect mental models or theories to pointing to instances where psychologists themselves ignore base rate information in interpreting their results.

Everyone can generate topics that should have been, but were not, included. I would have liked to see much more from the areas of social cognition and cross-cultural studies of thinking (to be fair, neither topic was ignored).

The way in which the book is perhaps most successful is in conveying a point of view. The search-inference framework does work for both descriptive analyses and the prescriptive conclusions. What comes across is not simply a textbook, but rather a person, Jonathon Baron, who is thinking and deciding about thinking and deciding. And following his own advice, the author almost always provides a balanced treatment of his material.

My only strong reaction within the framework of the author's goals concerns the sections on moral thinking. He argues for an universal rather than relativistic view of moral judgements and for a form of utilitarianism that defines best actions as those that have the highest expected utility across people. He says, 'If it is wrong for me to steal a book from the library, then it is wrong for you too, if you are in the same situation.' I have difficulty seeing how universalism follows from the rest of Baron's framework. If 'situation' includes personal goals and utilities, then it may be very unlikely that two people are in the same situation. Furthermore, who can determine what the utilities are? It goes against the spirit of the book to argue that society must impose utilities but if individuals are allowed to supply the utilities, then there will be utility in exaggerating one's own needs, wants and goals.

Maximizing expected utility across all people is just one of many options and a potentially dangerous one at that. Richard Atkinson (*American Psychologist*, 1972, 27, 921-931) provides an insightful discussion of alternative to maximizing expected utilities in employing computer-aided instruction effectively. He notes that maximizing the mean performance over all students may have the undesirable consequence of putting disadvantaged children still further behind the more advantaged children. Alternatively, one might aim to minimize the variance in performance or to maximize the mean performance subject to the constraint that the variance not be increased. I think maximizing average utility also carries with it the possibility for abusing minorities. If it's wrong for one person to jeopardize my health by smoking in my office, then it should not be less wrong for two people to smoke in my office (of course, there are other options, but I'm objecting to the logic that ties morality so closely to consensus). It is perhaps a minor quibble but at least

one piece of advice offered by the author cannot be computed. He suggests that the ideal behavior would provide 'the greatest benefit to others in return for the least sacrifice to ourselves.' The problem is that one cannot maximize these two functions simultaneously and it would seem to require a separate set of principles to determine how to weight them.

Despite these qualms about Baron's prescriptive analysis of moral judgement, his approach does have the virtue of consistency. I would much rather have a framework to agree or disagree with than a catalog of observations devoid of structure. Having said that, I wish to turn attention to the framework itself.

Although the framework that the author uses is successful, I find it somewhat incomplete. What is missing or perhaps understated is the analysis of the nature and complexity of the task that people face when they are thinking and deciding. Frequently, there is an unlimited set of information that one could potentially search and evaluate. As a result, extensive search is an impracticality. The computational complexity problem can be nicely illustrated by an example from Christopher Cherniak's (1986) book, *Minimal Rationality*. Suppose we have a computer that could evaluate the consistency of an assertion in the time that it takes a light ray to traverse the diameter of a proton, and suppose the computer ran for 20 billion years. Even if one had a procedure for doing this it would involve so many checks that there still would not be sufficient time to evaluate the internal consistency of a belief system containing only 138 logically independent propositions!

Given that many decision tasks are unconstrained with respect to potentially relevant information, it should not be surprising that people often adopt shortcut strategies or heuristics that are not optimal but which work pretty well most of the time (Herbert Simon refers to this as 'satisficing' rather than maximizing). Robin Hogarth (*Psychological Bulletin*, 1981, 90, 197-217) has argued that whether or not a strategy is functional depends on the conditions under which judgements are made. Much of the work on heuristics uses discrete judgement tasks. However, heuristics that are dysfunctional in a discrete task may be effective in continuous, interactive judgement tasks. Most of the time, Hogarth argues, it is good enough to get things going in the right direction with fine tuning being left for the future.

I would have liked to see a more extensive analysis of the costs associated with different decision procedures. For example, Payne, Bettman, and Johnson (*Journal of Experimental Psychology: Learning, Memory, and Cognition*, 1988, 14, 534-552) have shown through computer simulation that under time constraints, several (nonnormative) heuristics are more accurate (and therefore more adaptive) than a truncated normative procedure. They also observed that people adjusted design strategies in response to time

pressure in an adaptive manner (where adaptiveness is as again defined by computer simulations of the effectiveness of different strategies under different time costs).

Thinking and Deciding focuses on methods of thinking and mentions content effects only in passing. Again it's a matter of emphasis, but one wonders to what extent poor thinking can be attributed to knowledge rather than procedure. Mental models may be logically constructed but employ inadequate knowledge, as in the case of the woman who reasoned that her uterus was like a container that opened once a month during menstruation and that she would not become pregnant if she had intercourse only midway between her periods (when the container was safely closed).

Whether or not Johnson Baron's framework is incomplete, it does succeed in raising a fundamental question. Can we develop a coherent notion of good and poor thinking that is not so constrained that it imposes unrealistic expectations (thus overestimating the extent of poor thinking) nor so unconstrained that it becomes a vacuous standard (thus defining poor thinking out of existence)? *Thinking and Deciding* is a provocative attempt (in the best sense of the word) at addressing this issue.

DOUGLAS L. MEDIN
Psychology Department,
University of Michigan, Ann Arbor,
Michigan, USA

DECISION MAKING: DESCRIPTIVE, NORMATIVE AND PRESCRIPTIVE INTERACTIONS. David E. Bell, Howard Raiffa and Amos Tversky (Eds.). New York: Cambridge University Press, 1988. Cloth, 623 pages, index.

In June 1983 a conference was held at the Harvard Business School honoring the School's 75th anniversary with a gathering of illustrious decision theorists. The timing was splendid, as the editors note in their Preface to the current volume: '... not since the early 1950s ... has so much intellectual enthusiasm been directed at the question of how people should, and do, behave when called upon to take action in the face of uncertainty'. Given the setting, the occasion, and the distinction of the invited participants, one imagines the conference as a feast for the mind. For those unable to attend, perhaps the published record will recapture a little of the excitement.

Unfortunately, despite the excellence of many of the individual papers reprinted in this volume, not very much of the conference spirit survives. First, the editors' hope of including edited transcripts of the discussion tapes encountered '... technical and ... administrative difficulties' (p.5) and the transcripts were abandoned except for a dozen pages from the medical decision making session, reprinted here. Second, a good deal has happened in the area since 1983, and many of the papers have lost some of their 1983 freshness, having entered the literature in alternative forms. Third, a substantial fraction of the papers were not fresh-minted, even in 1983: of 28 papers included here, 11 have open-literature publication dates of 1983 or earlier; two more had already been issued as HBS Working papers by 1980; and seven more reached the open literature between 1984 and 1986. This leaves only eight pieces nominally 'new' for this volume, including the editors' introductory overview.

The volume, then, falls somewhat short as hot news from the front.

Where it works surprisingly well, in contrast, is a straightforward collection of excellent — in several cases classic — papers. We may not learn much about how the conference went, but we cannot miss the brilliance of the cast the organizers assembled, and of the papers each chose as a record of his or her participation. Indeed, if one considers the book not as a record of a conference but as a candidate text for a graduate seminar, it would score very highly, as a supplemental readings collection if not as the main text.

To lend some degree of orderliness to the collection, the editors have imposed a four-part structure. The first, entitled 'Conceptions of Choice', groups together nine quite heterogeneous pieces ranging from seminal theoretical contributions by Simon, March, Shafer and Fishburn to descriptive reviews by Einhorn and Hogarth, and Slovic, Fischhoff and Lichtenstein. A second, surprisingly brief, section includes four papers touching on aspects of probability in choice. A third section offers seven interestingly varied papers on utilities and values, including a delightful short paper from Schelling ('The Mind as a Consuming Organ'). The final section claims to represent applications, but does so only in a distant way except for three papers on medical decision making by Pauker, McNeil and colleagues. Four other pieces touch on applications in economic policy, business organizations, and education. The editors' introduction, in addition to a brave effort at covering this riot of diversity in a respectable mantle of orderliness, makes the case for adding a third category, prescriptive models, to the usual distinction between descriptive and normative. The case, though only modestly supported in the subsequent papers, is a strong one, if only as a reminder to those who would leap too swiftly from normative considerations to action recommendations that the gap to be closed is large.