PSYCHOLOGY REVIVALS

Cognitive Processes in Choice and Decision Behavior

Edited by
Thomas S. Wallsten

Psychology Revivals

Cognitive Processes in Choice and Decision Behavior

Decision theory is a uniquely interdisciplinary field of study with contributions from economics, statistics, mathematics, philosophy, operations research, and psychology. The 1970s had seen important changes in research on behavioral decision theory in terms of a shift from a reliance on economic and statistical models to an emphasis on concepts drawn from cognitive psychology. Originally published in 1980, *Cognitive Processes in Choice and Decision Behavior* contains papers that explore the reasons why these changes had come about and discuss the future directions to which they pointed. It was clear at the time that research in behavioral decision theory was changing dramatically. The chapters in this book represent a good assessment of the reasons the changes were coming about and some of the merits and problems of the directions in which it was moving. Today it can be read in its historical context.

Cognitive Processes in Choice and Decision Behavior

Edited by

Thomas S. Wallsten

Routledge
Taylor & Francis Group
LONDON AND NEW YORK

First published in 1980
by Lawrence Erlbaum Associates, Inc.

This edition first published in 2024 by Routledge
4 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

and by Routledge
605 Third Avenue, New York, NY 10017

Routledge is an imprint of the Taylor & Francis Group, an informa business

Publisher's Note
The publisher has gone to great lengths to ensure the quality of this reprint but points out that some imperfections in the original copies may be apparent.

Disclaimer
The publisher has made every effort to trace copyright holders and welcomes correspondence from those they have been unable to contact.

A Library of Congress record exists under ISBN: 089859054X

ISBN: 978-1-032-74505-3 (hbk)
ISBN: 978-1-003-46954-4 (ebk)
ISBN: 978-1-032-74508-4 (pbk)

Book DOI 10.4324/9781003469544

COGNITIVE PROCESSES IN CHOICE AND DECISION BEHAVIOR

edited by
THOMAS S. WALLSTEN
University of North Carolina at Chapel Hill

LAWRENCE ERLBAUM ASSOCIATES, PUBLISHERS
1980 Hillsdale, New Jersey

This work relates to Department of the Navy
Contract N00014-78-C-0170 issued by the Office
of Naval Research. The United States Government
has a royalty-free license throughout the world
in all copyrightable material contained herein.

Lawrence Erlbaum Associates, Inc., Publishers
365 Broadway
Hillsdale, New Jersey 07642

Library of Congress Cataloging in Publication Data

Main entry under title:

Cognitive processes in choice and decision
behavior.

Proceedings of a conference held June 22–24,
1978 at the University of North Carolina at
Chapel Hill.
"Contract N00014-78-C-0170 issued by the Office
of Naval Research."
Bibliography: p.
Includes index.
1. Decision-making—Congresses. 2. Choice
(Psychology))—Congresses. 3. Cognition—
Congresses. I. Wallsten, Thomas S.
BF441.C53 153.8'3 79-27553
ISBN 0-89859-054-X
Printed in the United States of America

Contents

Preface

Decision theory is a uniquely interdisciplinary field of study with contributions from economics, statistics, mathematics, philosophy, operations research, and psychology. Recent years have seen important changes in research on behavioral decision theory in terms of a shift from a reliance on economic and statistical models to an emphasis on concepts drawn from cognitive psychology. In order to explore the reasons why these changes have come about, and to discuss the future directions to which they point, a conference was held from June 22–24, 1978, at Quail Roost, an idyllic conference center run primarily for the University of North Carolina at Chapel Hill. This volume contains the proceedings of that conference, and should be of interest to cognitive psychologists, decision theorists, decision analysts, and related scientists.[1]

The schism that, until recently, has existed between behavioral decision theory and the rest of cognitive psychology has been unfortunate, although understandable. Pitz (1977) is probably correct in his analysis that this separation occurred because the the roots of decision theory lie squarely in economics and statistics, whereas those of cognitive psychology can be found in the early schools of association and rationalism. Estes (this volume) presents a similar perspective. It is useful to consider certain aspects of decision-oriented and cognitive research to see how each can benefit from the other and to understand why the two fields may be growing closer together.

Research in behavioral decision theory has been concerned primarily with developing, testing, and reformulating relatively sophisticated formal models,

[1]A paper by Kahneman and Tversky (1979) was also presented at the conference, but is not included here because it has been published elsewhere.

most of which are normative in character. This work has focused both on global evaluations of the models and on testing various axioms from which the models flow, on measuring subjective probability and utility, and on developing probalistic models to pit against the normative algebraic ones. (Reviews of much of this earlier literature can be found in Becker & McClintock, 1967; Edwards, 1954, 1961; and Rapoport & Wallsten, 1972.)

Although individual models have been successful in the sense that a particular model can account for a good deal of the variance in a particular situation, there is a feeling among many researchers that the overall approach has not been fruitful. For one thing, it has not been possible to apply the results from one paradigm or with respect to one model to other paradigms or other models in any satisfying way. Perhaps more important than the lack of generalizability have been the findings that various axioms are systematically violated under a variety of conditions, utility is frequently not risk invariant, and people often do not make decisions so as to optimize some well-specified objective function. This apparent lack of progress in psychological decision theory was made starkly evident in the preface to *Contemporary Developments in Mathematical Psychology* (Krantz, Atkinson, Luce, & Suppes, 1974) in which the editors wrote:

> Perhaps the most striking exclusion (in the set of topics covered) is the entire area of preferential choice. There is no lack whatever of technically excellent papers in this area, but they give no sense of any real cumulation of knowledge. What are the established laws of preferential choice behavior? . . . [p. *xii*].

It is decidely not the case that researchers have simply catalogued descriptive successes or failures of normative models. Rather, in the search for more complete or useful descriptive models, investigators have been led to concepts and findings in various areas of cognitive psychology (see Hogarth, 1975, or Slovic, Fischhoff, & Lichtenstein, 1977). The alternative theories being suggested derive from the acknowledgment that decision making is a complex cognitive task, frequently situation dependent, which humans perform in a manner determined by their limited memory, retention, and information-processing capabilities. In certain respects, recent developments are similar to those advocated by Simon (1957) and, as Lockhead (this volume) points out, also by Bruner, Goodnow, and Austin (1956). Discussions of how cognitive limitations affect decision processes appear repeatedly throughout this book, but can be found explicitly in the chapters by Einhorn; Pitz; Payne; Fischhoff, Slovic, and Lichtenstein; Lockhead; MacCrimmon, Stanbury, and Wehrung; Wallsten; and Estes.

It is clear that decision researchers have come to realize the importance of cognitive concepts and cognitively oriented theories in understanding choice behavior. However, it is less obvious that cognitive psychologists yet acknowledge the importance of choice behavior in understanding intellectual processes such as memory, problem solving, letter recognition, or the like. This point is developed by both Estes and Lockhead in this volume. Indeed, because the majority of tasks studied by cognitive psychologists involve people making choices of one sort or another, Lockhead goes so far as to suggest that perhaps we should study "choice and decision behavior in cognitive processing" rather than the reverse.

Both Kubovy and Healy (this volume) and Estes (this volume) remind us that signal-detection theory is a normative choice model employed in a wide range of cognitive theories. Furthermore, the distinction between amount of information and criterion for a choice implied by this theory is also implicitly accepted by many cognitive theorists. However, in general, this decision aspect of a cognitive theory is relegated to a black box insensitive to the context within which it is placed, and rarely is the area of behavioral decision theory called upon to supply helpful concepts or findings for the purpose of improving the cognitive theory. It is to be hoped that these proceedings might stimulate cognitive psychologists to attend more carefully to the decision aspects of their subjects' tasks.

Both Lockhead and Estes point out that, generally, decision theorists have been concerned with formal descriptions of the task environment and with optimal strategies for performing such tasks, and consequently, have studied behavior within a narrowly defined range of situations. Alternatively, cognitive theorists have been concerned primarily with processes that cut across tasks and consequently have studied behavior in a wider range of richer but less well understood environments. The chapters in this book represent clear attempts to merge these two approaches.

Some of the participants in this conference were invited to describe the present structure of their theoretical framework, indicating the roots from which it grew, how it ties in with other areas of psychology or decision theory, and future prospects stemming from it. It was anticipated that these papers would fall relatively neatly into certain classifications that would form a basis for partitioning this book. Other participants (Carroll, Estes, Lockhead, and Luce) were each invited to discuss and provide commentary on a specific set of papers from the perspective of his particular specialization in psychology. The papers that were assigned to the discussants, naturally enough, are those that immediately precede their chapters in the book. All the authors responded to their invitations in such a comprehensive fashion that, happily, it has become impossible to classify the papers. There are numerous interrelated messages in each paper, and furthermore, the discussion chapters for the most part extend so far beyond the specific papers assigned to them that they stand as useful and important contributions in their own right.

Broadly speaking, three themes are woven throughout all the chapters. One is that we must enlarge the range of paradigms studied. Another is that we must broaden the scope of the underlying psychological theories employed. The third is that we must utilize mathematical models in less simplistic fashions. Truthfully and trivially, charges of this sort can be leveled against all research in all areas.

The contributions of the chapters lie not in the charges, but in the directions of the solutions they propose.

If one wishes to read chapters that focus to a substantial degree on fruitful ways we can enlarge the range of paradigms, then one would turn to Ebbesen and Konečni; Corbin; Payne; Carroll; Fischhoff, Slovic, and Lichtenstein; Lockhead; Schum; and Estes. Ebbesen and Konečni discuss their recent research relating laboratory and field studies of legal decision making. They find systematic differences in behavior between college students and legal professionals in the laboratory, and also between the behavior of the professionals in the laboratory and in the real world. They go on to suggest that although one might use laboratory experiments to study specific cognitive limitations, or processes, it would be a mistake to imagine that there exist a small number of laws of decision behavior that can be uncovered in the laboratory and then applied in a straightforward way to real world decisions. Thus, laboratory and observational studies should proceed hand-in-hand.

Corbin, in her chapter, suggests that we will learn much more about decision processes by studying the determinants that inhibit or allow choices to be made than by studying the choices themselves. Thus, we should focus to a considerably larger extent than we have on prechoice environments and behaviors. By providing a conceptual organization to the range of barriers that must be overcome prior to the making of a choice, Corbin suggests a framework for future empirical and theoretical research.

Payne, too, provides ways to usefully study the decision process from the subject's first introduction to the task, through his or her understanding of it, to the final set of choices. Payne suggests that a range of measures, such as verbal protocols, order of information search, eye movements, and so on, be collected to provide a fuller understanding of the process.

Carroll suggests that our models have dictated to too large an extent the paradigms in which we have collected data. If we consider decision makers as adaptive and at the same time beset with certain cognitive limitations, we will realize that they are adapting to the task as they view it, and consequently that we must expand the range of situations studied.

Fischhoff, Slovic, and Lichtenstein are concerned primarily with the elicitation of value judgments from decision makers. They point out that often these values are poorly defined or formulated by the subject, and that, as a result, the particular judgments elicited will depend to a large degree on the method of questioning. These authors demonstrate that we may achieve greater insight into the nature of people's values by posing diverse questions and studying the nature of the apparently inconsistent responses.

Lockhead shows the close theoretical correspondence among paradigms studied in decision research, problem solving, and certain aspects of psychophysics. The specific questions asked in each of these areas are relevant to the other areas, and Lockhead demonstrates ways in which it would be beneficial to look across the paradigms.

Schum is interested in the inductive use of equivocal information when that information is nonindependent and hierarchically related to the hypotheses in question. In itself, this is an important advance over the usual paradigms involving simple probabilistic linkages between hypotheses and data. However, Schum goes on to show that formal models and behavioral theories of the process can be aided by studying the law concerning use, interpretation, and admissibility of courtroom evidence.

Estes' chapter relates contemporary behavioral decision research to certain long-term trends in psychology, and in so doing suggests a variety of ways that our paradigms might be broadened. We should, for example, devote more attention to experimental situations in which the choice alternatives are not well defined, in which memory for information can be assessed, and in which individual differences can be systematically explored.

If one wishes to read chapters emphasizing ways in which psychological aspects of decision theories can be enriched, then one would look at those by Einhorn; Pitz; Lockhead; MacCrimmon, Stanbury, and Wehrung; Wallsten; Kubovy and Healy; and Estes.

Einhorn demonstrates that theories of learning must be included in our understanding of the decision process. People learn action-outcome linkages, and frequently they see causal linkages where none exist. The question is, how do people's experiences give rise to the range of normative and heuristic rules that they bring to bear in various situations? Einhorn suggests that we must particularly study the nature of outcome feedback, reviews his research on that topic, and develops a theory concerning how subjects learn in a choice situation based on their misinterpretation of feedback.

Pitz relies heavily on Newell and Simon's (1972) "production systems" to develop a class of theories concerning how people encode and process the distributional properties of outcomes. Within this context, he demonstrates how subjects' internal representations of tasks can depend on certain cognitive limitations, and how heuristic rules can be derived.

Payne also relies heavily on the information processing theory of Newell and Simon to suggest that we develop models that include the subject's internal representation of the environment, or his or her problem space, as well as descriptions of the environment itself. Payne reviews his research showing that the subject's problem space, and therefore his or her decision strategy, depends on the task and on how it is presented.

As previously indicated, Lockhead argues that decision processes, problem solving, and psychophysics can profitably be studied jointly in a manner that will enhance the commonalities among the theories involved.

MacCrimmon, Stanbury, and Wehrung employ a very simple mathematical model of risk to analyze the data obtained from their business executive subjects, and as a consequence demonstrate the importance of context on people's choices. They interpret these data by developing a theory incorporating selective perception and simple decision making.

Wallsten's chapter presents a general theory relating selective attention and simple task-specific decision rules to a wide range of choice situations. The theory is developed in a manner intended to be consistent with the findings on bounded rationality and heuristics, but is formulated so as to allow specific predictions and the relation of behavior in one situation to that in another. The approach is illustrated by applying it to the study of probabilistic inference.

Within the framework of signal-detection theory, Kubovy and Healy present and study classes of psychological theories concerning probabilistic inference, or categorization, as they call it. They are concerned in particular with how subjects learn, form their choice rules and decision criteria, and evaluate the probabilistic nature of information. Kubovy and Healy discuss some of their research that rules out, or makes less likely, certain classes of theories, each of which encompasses various specific models.

Finally, Estes suggests ways in which decision theorists might focus less strongly on particular tasks and devote more attention to developing theories about basic processes that cut across tasks. He proposes that the distinction currently made in many areas of cognitive psychology between structural and control processes will be useful in understanding decision behavior, and in relating decision processes to other areas of psychology.

The chapters that provide explicit examples of how mathematical models can be applied to behavioral decision research in less simplistic ways include those by MacCrimmon, Stanbury, and Wehrung; Schum; Wallsten; and Kubovy and Healy.

As already mentioned, the analyses by MacCrimmon, Stanbury, and Wehrung are guided by relatively elementary mathematical models of risk. However, the relationship between model successes or failures and features of the choice alternatives is traced in such a fashion that our knowledge of the determinants of subjective risk is enhanced considerably.

Schum specifies classes of Bayesian models that tease out and formalize the logical connections between evidence and hypotheses when the evidence is indirect, nonindependent, and hierarchical. Such a situation occurs, for example, when multiple unreliable witnesses report an event. This allows Schum to develop a scheme for classifying evidence in terms of its source and the nature of its relationship to the facts at issue. Schum's approach provides a framework for systematically studying complex inferences and relating it to other cognitive processes. In his discussion of Schum's chapter, Luce uses concepts from signal-detection theory to demonstrate some of the problems involved in combining multiple reports of an event.

Wallsten's chapter makes use of algebraic composition rules in a manner that explicitly interprets the parameters in terms of psychological constructs. This provides a means for relating predictive failures of models to substantive theory and for generalizing results from one paradigm to another. The focus of research thus shifts from whether a model is right or wrong to the development of a general descriptive theory of the decision process, which, however, is modeled differentially, depending on the task.

Kubovy and Healy superimpose on signal-detection theory formal representations of learning processes and of various psychological considerations that could lead to suboptimal performance. This approach provides a taxonomy of theories for the probabilistic categorization task and a systematic means for evaluating the theories.

It is clear that research in behavioral decision theory is changing dramatically. The chapters in this book represent a good assessment of the reasons the changes are coming about and some of the merits and problems of the directions in which we are moving. In that sense, the chapters are speculative, and as such, there is more than occasional disagreement between them. I hope the result is thought provoking to the reader.

I express my sincere appreciation to the authors for their thorough, thoughtful, and timely responses to my editorial comments. Their cooperation made my job as editor far easier and more enjoyable than I was led to believe it would be. Special thanks are due Michael Kubovy for assistance in organizing the conference, Curtis Barton for handling many of the details, and Elizabeth Schopler for secretarial assistance above and beyond the call of duty. The conference was made possible by support from the Office of Naval Research through contract N00014-78-C-0170.

Thomas S. Wallsten

Chapel Hill, North Carolina

REFERENCES

Becker, G. B., & McClintock, C. G. Value: Behavioral decision theory. *Annual Review of Psychology,* 1967, *18,* 239-286.

Bruner, J. S., Goodnow, J. J., & Austin, G. A. *A study of thinking.* New York: Wiley, 1956.

Edwards, W. The theory of decision making. *Psychological Bulletin,* 1954, *51,* 370-417.

Edwards, W. Behavioral decision theory. *Annual Review of Psychology,* 1961, *12,* 473-498.

Hogarth, R. Cognitive processes and the assessment of subjective probability distributions. *Journal of the American Statistical Association,* 1975, *70,* 271-293.

Kahneman, D., & Tversky, A. Prospect theory: An analysis of decision and risk. *Econometrica,* 1979, *47,* 263-291.

Krantz, D. H., Atkinson, R. C., Luce, R. D., & Suppes, P. (Eds.), *Contemporary developments in mathematical psychology* (Vol. 1). San Francisco: Freeman, 1974.

Newell, A., & Simon, H. A. *Human problem solving.* Englewood Cliffs, N.J.: Prentice-Hall, 1972.

Pitz, G. F. Decision making and cognition. In H. Jungermann & G. De Zeeuw (Eds.), *Decision making and change in human affairs*. Dordrecht, Holland: Reidel, 1977.

Rapoport, A., & Wallsten, T. S. Individual decision behavior. *Annual Review of Psychology,* 1972, *23,* 131–176.

Simon, H. A. *Models of man: Social and national.* New York: Wiley, 1957.

Slovic, P., Fischhoff, B., & Lichtenstein, S. Behavioral decision theory. *Annual Review of Psychology,* 1977, *28,* 1–39.

1 Learning from Experience and Suboptimal Rules in Decision Making

Hillel J. Einhorn
Graduate School of Business
Center for Decision Research
University of Chicago

Current work in decision-making research has clearly shifted from representing choice processes via normative models (and modifications thereof) to an emphasis on heuristic processes developed within the general framework of cognitive psychology and theories of information processing (Payne, this volume; Russo, 1977; Simon, 1978; Slovic, Fischhoff, & Lichtenstein, 1977; Tversky & Kahneman, 1974, 1980). The shift in emphasis from questions about how well people perform to how they perform is certainly important (e.g., Hogarth, 1975). However, the usefulness of studying both questions together is nowhere more evident than in the study of heuristic rules and strategies. This is because the comparison of heuristic and normative rules allows one to examine discrepancies between actual and optimal behavior which then raise questions regarding why such discrepancies exist. The approach taken here is to focus on how one learns both types of rules from experience. The concern with learning from experience raises a number of issues that have not been adequately addressed; e.g., under what conditions are heuristics learned? How are they tested and maintained in the face of experience? Under what conditions do we fail to learn about the biases and mistakes that can result from their use?

The importance of learning for understanding heuristics and choice behavior can be seen by considering the following:

1. The ability to predict when a particular rule will be employed is currently inadequate (Wallsten, this volume). However, concern for how and under what conditions a rule is learned should increase one's ability to predict when it is likely to be used. For example, if a rule is learned in situations in which there is

DOI: 10.4324/9781003469544-1

little time to make a choice, prediction of the use of such a rule is enhanced by knowing the time pressure involved in the task.

2. A concomitant of (1) is that it should be possible to influence how people judge and decide by designing situations in which tasks incorporate or mimic initial learning conditions. The implications of this for both helping *and* manipulating people are enormous (Fischhoff, Slovic & Lichtenstein, 1978; this volume).

3. Consideration of learning focuses attention on environmental variables and task structure. Therefore, variables such as amount of reinforcement, schedules of reinforcement, number of trials (= amount of experience), and so on, should be considered in understanding judgment and decision behavior (cf. Estes, 1976). Although the importance of the task for understanding behavior has been continually stressed (Brunswik, 1943; Castellan, 1977; Cronbach, 1975; Dawes, 1975; Edwards, 1971; Einhorn & Hogarth, 1978; Simon & Newell, 1971), psychologists seem as prone to what Ross (1977) calls the "fundamental attribution error" (underweighting environmental factors in attributing causes) as anyone else.

4. A major variable in understanding heuristics is outcome feedback. Because outcome feedback is the main source of information for evaluating the quality of our decision/judgment rules, knowledge of how task variables both affect outcomes and influence the way outcomes are coded and stored in memory becomes critical in explaining how heuristics are learned and used.

5. The area of learning is the focal point for considering the relative merits of psychological versus economic explanations of choice behavior. Some economists have argued that although one does not act "rationally" all the time, one will learn the optimal rule through interaction with the environment. Vague assertions about equilibrium, efficiency, and evolutionary concepts are advanced to bolster this argument. Therefore, study of how (and how well) people learn from experience is important in casting light on the relative merits of psychological and economic theories of choice.

LEARNING FROM EXPERIENCE: HOW?

It is obvious that decision making is action oriented; one has to choose what action to take in order to satisfy basic needs and wants. Therefore, it is important for any organism to learn the degree to which actions will lead to desirable or undesirable outcomes. This means that a great deal of learning from experience must involve the learning of action-outcome linkages. Furthermore, because actions and outcomes are contiguous, people are prone to interpret the links between them as representing cause-and-effect relationships (Michotte, 1963). Therefore, the strong tendency to see causal relations can be considered an outgrowth of the need to take action to satisfy basic needs. Moreover, as pointed out by Tversky and Kahneman (1980), the learning of causal relationships and the organizing of events into causal "schemata" allow people to achieve a coherent interpretation of their experience. Finally, the learning of action-outcome links is important for understanding how people learn their own tastes or utilities. For example, consider a child who chooses a particular vegetable to eat, experiences an unpleasant taste, and thereby learns to associate a negative utility with that food. Note that it is typically by choosing that consequences can be experienced and utility learned. Therefore, the learning of action-outcome links and the learning of utility are closely tied together.

Although we learn from experience by taking action, how does one initially learn which alternative to choose? Undoubtedly, much initial learning occurs by trial and error—that is, people randomly choose an option and observe the outcome (cf. Campbell, 1960). The process by which trial-and-error learning gives way to the development of strategies or rules is not well known (cf. Siegler, 1978). However, one can speculate that both reinforcement from trial-and-error learning and generalization (both stimulus and response) play an important role (Staddon & Simmelhag, 1971). In any event, the rules we develop seem directly tied to learning what outcomes will follow from particular actions. As previously described, learning from experience is basically inductive in nature; that is, one experiences specific instances or cases and heuristics are developed to provide some generality for dealing with them. The inductive nature of learning from experience has several implications regarding heuristics:

1. Specificity of Rules. If learning occurs inductively via specific cases, then heuristic rules should be extremely context dependent. Much evidence now suggests that this is indeed the case (Grether & Plott, 1979; Lichtenstein & Slovic, 1971; Simon & Hayes, 1976; Tversky & Kahneman, 1980). The way in which a problem is worded or displayed, or a particular response is asked for, all seem to make an important difference in the way information is processed and responses generated. A dramatic example of this specificity can be seen in the work of Simon and Hayes (1976) on "problem isomorphs." They have shown that different surface wordings of structurally identical problems (i.e., problems that can be solved using identical principles) greatly change how people represent the problem in memory and consequently solve it. An important implication of this result is that in order to make heuristic models more predictive, one must contend with the task as represented and not necessarily with the task structure as seen by an experimenter. A particularly timely example of the importance of this phenomenon in predicting behavior is provided by observing that behavior depends on whether a tax cut is represented as a gain or a smaller loss (Kahneman & Tversky, 1979).

2. Generality of Rules. If heuristics are rules learned through induction, it is necessary to group tasks by similarity or else there would be as many rules as situations. Because this latter possibility is unacceptable, heuristics must have some generality over tasks. However, this conclusion contradicts what was said previously about context dependence and specificity of rules. This paradox can be resolved if one considers the range of tasks to which a rule can be applied. For example, consider the rule: "Never order fish in a meat restaurant." Although such a rule is general with respect to a certain type of restaurant, it is certainly more specific than the rule: "Judge the probability with which event B comes from process A by their degree of similarity" (Tversky & Kahneman, 1974). The latter heuristic is clearly at a much higher level of generality. In fact, it may be that heuristics such as representativeness, availability, anchoring and adjusting, are "metaheuristics"—that is, they are rules on how to generate rules. Therefore, when confronted by problems that one has not encountered before (such as judging probabilities of events), or problems whose specificity makes them seem novel, metaheuristics direct the way in which specific rules can be formed to solve the problem. The idea of a metaheuristic allows one to retain the generality that any rule necessarily implies, yet at the same time allows for the important effects of context, wording, response mode, and so on. In order to illustrate, consider the study by Slovic, Fischhoff, and Lichtenstein (1976) in which people were asked to judge the relative probabilities of death from unusual causes. For example, which has a higher probability: being killed by lightning or dying from emphysema? When confronted with such a question, there are many ways to attempt an answer. One rule that could be used would be: "Think of all the people I know that have died from the two causes and pick the event that caused more deaths." In my own case, I would choose emphysema (which does have a higher probability, although most people pick being killed by lightning). However, I could have just as easily developed a rule that would lead to the opposite answer; e.g., "Think of all of the cases of being killed by lightning and of death from emphysema that I have ever *heard about* (newspapers, television, etc.)." If this were my rule, I would choose being killed by lightning as being more probable. Note that in both cases I have used an availability heuristic. Clearly, the way in which a question is phrased could induce specific rules that lead to different results, yet these specific rules could be classified under a single more general strategy, or metaheuristic.

3. Strength of Heuristics. If heuristics are learned inductively, then learning occurs over many trials with many reinforcements. As is discussed later, because of the way feedback occurs and the methods that we use to test rules via experience, positive reinforcement can occur even for incorrect rules (Wason, 1960). Moreover, in addition to the large number of reinforcements that we experience, the size or intensity of reinforcement can be large. For example, gaining a sizable amount of money following the use of some rule for picking stocks should have a considerable reinforcement effect. Therefore, unlike laboratory studies of human learning, in which ethical considerations prevent large positive and negative reinforcements, our own experience poses no such constraints.

LEARNING FROM EXPERIENCE: HOW WELL?

The question of how well we learn from experience focuses attention on comparing heuristic rules to optimal rules. Therefore, it must be asked how the latter are learned and what the implications are for applying them in our own experience? Optimal rules, such as Bayes' theorem, optimization, and so on are learned *deductively*. In fact, much of what can be called formal learning is of a deductive character; that is, we are taught scientific laws, logical principles, mathematical and statistical rules, and so on. Such rules are by their very nature abstract and context independent. Furthermore, when context can influence the form of a rule, one is frequently told that the rule holds, "other things being equal." Of course, in our own experience, other things are rarely equal, which makes the learning of optimal rules via induction so difficult. (The original discoverers or inventors of optimal rules overcame these difficulties; however, this distinguishes them from the rest of us.)

The abstract nature of deductive rules has important implications regarding the difficulty people have in applying optimal methods to specific situations. This difficulty centers around the ability to discern the structure of tasks that are embedded in a rich variety of detail. Therefore, when one is faced with a specific problem that is rich in detail, and in which details may be irrelevant or redundant, one's attention to specifics is likely to divert attention from the general structure of the problem. In fact, the very abstractness of deductively learned optimal rules may prevent them from being retrieved from memory (cf. Nisbett, Borgida, Crandall, & Reed, 1976). Therefore, abstract rules may not be very "available" in specific cases. However, this begs the question because it is important to know *why* these rules are not available.

Consider the way action-outcome combinations are likely to be organized and stored in memory. In particular, consider whether such information is more likely to be organized and stored by content or task structure. It would seem easier and more natural to organize action-outcome combinations by subject matter rather than by structure; for example, experiences with schools, parents, members of the opposite sex, and so on, rather than Bayesian problems, selection situations, optimization problems, and the like. That content can differ while structure remains the same is quite difficult to see (Kahneman & Tversky, 1979; Simon & Hayes, 1976). Therefore, I think it unlikely that most people organize their experiences by task structure. This is not to say that one could not be trained to do so. In fact, much of professional training is exactly this; for instance, one is taught to recognize problems as belonging to a class of problems having a given structure and (sometimes) a known solution. Therefore, optimal rules can be "available" through extensive training. Of course, there is the danger of such rules being *too* readily available; that is, problems are forced into a structure that is not appropriate because a solution within that structure exists. It is a truism that when presented with a problem, professionals view the problem within the structures they have been trained to see. Therefore, although professional training does involve a concern for structure, such training is generally within a narrowly defined content area.

Further evidence illustrating the need to group problems by content rather than structure is provided by considering the way public knowledge about the world is organized and taught. For example, departmentalized education, professional training, cataloguing of information in libraries and encyclopedias, and so on, illustrate the organizing of information by content rather than structure. Although there are great advantages in organizing knowledge in this way, there are also costs. The difficulty of applying optimal rules developed in one content area to structurally similar problems in other content areas may be one such cost. However, at the level of the individual learner, other difficulties are now considered that may be even more costly.

Although task structure is difficult to discern, outcomes are not; they are highly visible, available, and often unambiguous. Therefore, consideration of reinforcement via outcome feedback is essential in understanding how heuristics are maintained in the face of experience. Furthermore, if outcomes are a function of task structure to a considerable degree and the decision maker's knowledge of such structure is lacking, then rules that are irrelevant or even poor may still be reinforced by positive outcome feedback (for example, "superstitious" behavior in animal learning; see Staddon & Simmelhag, 1971).

Following are two examples of how normatively poor heuristics can lead to good outcomes and in which awareness of the poor quality of the rule may be lacking. Consider shopping in a supermarket and coming to cans of juice with the following prices and overall quality levels (adapted from Tversky, 1969):

Brand	*Price*	*Quality*
X	40¢	High
Y	35¢	Medium
Z	30¢	Low

Assume that I use the following rule to choose between the three brands: If the price difference is 5¢ or less, choose the brand with the higher quality; if the price difference is greater than 5¢, choose according to price. Such a simple rule (which is a lexicographic semiorder) leads to:

$$X > Y$$
$$Y > Z, \text{ but}$$
$$Z > X$$

Therefore, this rule leads to intransitive choices, which are clearly irrational. However, note that after I choose X over Y, I may then eliminate Y from the remaining set and compare X with Z. Therefore, I end up with Z, which may be quite acceptable after I taste it. I then congratulate myself on what a good shopper I am—I saved money and I got a reasonable product. The important point to note here is that by not making the Y vs. Z comparison, I remain unaware that my rule leads to an intransitive choice. All I *am* aware of is that I made a choice with minimal fuss and strain and the outcome was satisfactory. Therefore, positive outcome feedback reinforces a normatively poor rule, and awareness that something is wrong is missing.

The second example is a probabilistic one (cf. Schum, this volume). Imagine that you are a military general in a politically tense area and that you are concerned that your enemies will invade your country. Furthermore, from past experience it is known that when enemy troops mass at a border, the probability of invasion is .75. However, you do not have direct access to information about enemy troops, but must rely on a report of such activity by your intelligence sources. As it turns out, every time your intelligence sources report that troops are massing, they are really there. Consider that you now receive a report from your sources that enemy troops are at the border. What is the probability of invasion? More formally, let:

H = hypothesis of being invaded
D = troops massing at the border
D^* = report of troops massing at the border

The problem states that $p(H|D) = .75$ and $p(D|D^*) = 1.0$, and asks you for $p(H|D^*)$. If you are like most people, you probably answered .75. However, the information given is not sufficient to answer the question in the normatively correct way. In fact, it is possible that in the preceding problem, $p(H|D^*) = 0$. Because most people find this very difficult to believe, consider Fig. 1.1, which illustrates the problem by means of a Venn diagram. Note that the intersection of H with D^* is null, so that the conditional probability, $p(H|D^*)$, is zero. The reason that people find this result so surprising is that they have made a logical

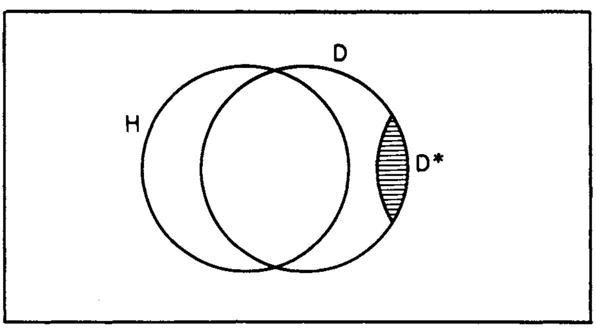

FIG. 1.1. Venn diagram showing the relationship between the hypothesis (H), datum (D), and report of datum (D^*).

fallacy of the form: If D^* implies D, then D implies D^*. Although D occurs whenever D^* is given, the reverse is not necessarily the case. In fact, an intuitive way to see the issue is to think that the enemy is particularly cunning so that your intelligence sources see their troops only when there is no invasion planned. However, when an invasion is planned and troops are at the border, they are hidden so that your sources do not report them.

This example illustrates the difficulty of applying optimal rules (in this case, the rules of formal logic) to a specific task. Although very few people would make the logical error when it is presented in a recognizable form, the importance of the example is that it shows how the specifics of the problem hide its real structure so that optimal rules are easily violated (cf. Kahneman & Tversky, 1979). A second point can be made with respect to this example. Consider that the general makes the logical error and estimates the chance of war at .75. He then sends *his* troops to the border, thereby causing an invasion by the enemy. Therefore, the faulty reasoning of the general is reinforced by outcome feedback—"after all," he might say, "those SOB's did invade us, which is what we thought they'd do."

The two examples just discussed illustrate the basic point of this chapter, viz. without knowledge of task structure, outcome feedback can be irrelevant or even harmful for correcting poor heuristics. Moreover, positive outcome feedback without task knowledge tends to keep us unaware that our rules are poor because there is very little motivation to question how successes were achieved. The conditions under which outcome feedback does not play a correcting role vis-a-vis heuristics and strategies are denoted *outcome irrelevant learning structures* (OILS). Such structures may be much more common than we think. Before examining one such structure in detail, consider probabilistic judgments within the framework of OILS, because much of the work on heuristics is directly concerned with this type of judgment. Consider that you judge the probability of some event to be .70. Let us say that the event does not happen. What does this outcome tell you about the quality of the rules used to generate the judgment? One might argue that any *single* outcome is irrelevant in assessing the "goodness" (i.e., degree of calibration) of probabilistic judgments. Therefore, in an important sense, immediate outcome information is irrelevant for correcting poor heuristics. It is only if one keeps a "box score" of the relative frequency of outcomes when one judges events with a given probability that one can get useful feedback from outcomes. However, this is likely to be a necessary but not sufficient condition for making well-calibrated judgments. First, over what time period does one keep the box score before deciding that the judgment is or is not calibrated? Furthermore, how close is close enough in order to say that the judgment is accurate (in the sense of being well calibrated)? Note that this whole mode of evaluating outcomes involves reinforcement that is delayed for long time periods. Therefore, it is not clear that such feedback will have much of a self-correcting effect. Second, in order to learn about the goodness of rules for estimating probability, one's box score must include not only one's estimates and the resulting outcomes, but also one's rules for deriving those estimates. For example, if I kept a record of outcomes that resulted for 100 cases in which I gave estimates of .70, what would the information that 53 of those times the event happened tell me about the quality of the rules I used? Because it is likely that many different rules could have been used to estimate probabilities in the 100 different situations, the outcome information is irrelevant and outcome feedback is not useful unless one is both aware of one's rules and a record is kept of their use (cf. Nisbett & Wilson, 1977 on whether we are aware of our own cognitive processes).

The preceding example does not imply that it is impossible to learn to make well-calibrated probability judgments. If one makes *many* probability judgments in the *same situation*, such as weather forecasters and horse-racing handicappers do, and outcome feedback is quickly received, such conditions may not be outcome irrelevant, and feedback can be self correcting. However, such conditions would seem to be the exception rather than the rule for most of us.

Although probabilistic judgments typically occur in OILS, what about nonprobabilistic judgments? Surely, if one makes a prediction about something, one can check to see if the prediction is correct or not. Therefore, it would seem that outcomes should be relevant for providing self-correcting feedback. The remainder of this chapter discusses this issue within the context of one general and prevalent task structure, although the specific content of such tasks may be quite different.

SELECTION TASK[1]

A very general task involving nonprobabilistic judgments is now examined because outcome information seems both available and relevant for providing self-correcting feedback. The task considered is one in which judgments are made for the purpose of choosing between alternative actions. For example, consider a situation with two possible actions, A and B. Denote by x an overall, evaluative judgment, which may itself be a function of various types and amounts of information. Furthermore, let x_c be a cutoff point such that:

$$\begin{aligned}&\text{if } x \geqslant x_c, \text{ take action } A;\\ &\text{if } x < x_c, \text{ take action } B.\end{aligned} \tag{1.1}$$

Although simplistic, Eq. (1.1) applies to many judgment/decision situations, for example: job hiring, promotion, admission to school, loan and credit granting, assignment to remedial programs, admission to social programs, journal article acceptance, grant awarding, and so on. In these cases, a judgment of the degree

[1]Much of this section is drawn from Einhorn & Hogarth (1978).

of "deservedness" typically determines which action is to be taken because the preferred action cannot be given to all.

In order to compare judgment to a standard, the existence of a criterion, denoted y, is assumed to serve as the basis for evaluating the accuracy of judgment. Although the practical difficulties of finding and developing adequate criteria are enormous, the focus here is theoretical: It is the concept of a criterion that is necessary for this analysis. To be consistent with the formulation of judgment, it is further assumed that the criterion has a cutoff point (y_c) such that $y \geqslant y_c$ and $y < y_c$ serve as the basis for evaluating the outcomes of judgment. Thus, as far as learning about judgment is concerned, representation of outcomes in memory is often of a categorical form, i.e., successes and failures (cf. Estes, 1976).

It is very important to note that the structure of the task is one in which judgments (predictions) lead to differential actions and that outcomes are then used as feedback for determining the accuracy of the predictions. The formal structure can be seen by considering the regression of y on x and the four quadrants that result from the intersection of x_c and y_c, as illustrated in Fig. 1.2.

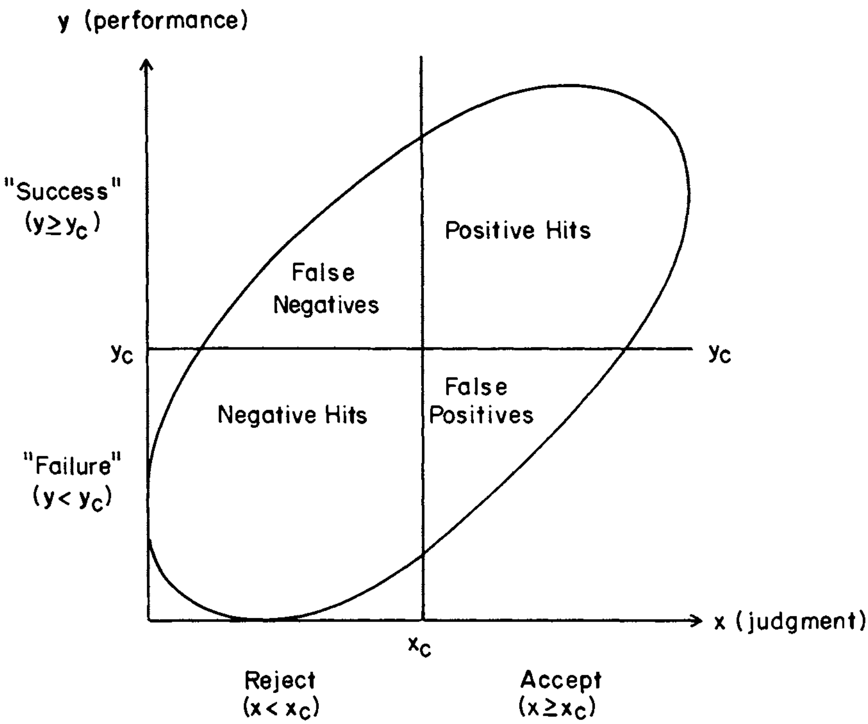

FIG. 1.2. Action-outcome combinations that result from using judgment to make an accept/reject decision.

Denote the correct predictions as positive and negative hits, and the two types of errors as false positives ($y < y_c|x \geqslant x_c$) and false negatives ($y \geqslant y_c|x < x_c$). To estimate the relationship between x and y (i.e., the correlation between x and y, ρ_{xy}) it is necessary to have information on *each* judgment-outcome combination. Assume first that such information becomes available over time (i.e., sequentially) and consider the experimental evidence concerned with learning the relationship between x and y in such circumstances. Research on the ability to judge the contingency between x and y from information in 2×2 tables (Jenkins & Ward, 1965; Smedslund, 1963, 1966; Ward & Jenkins, 1965) indicates that people judge the strength of relationship by the frequency of positive hits (in the terminology of Fig. 1.2), while generally ignoring information in the three other cells. These results are extremely important because they indicate that *even when* all of the relevant outcome information is available, people do not use it. This means that in laboratory studies that have outcome-relevant learning structures, people have transformed them into outcome-irrelevant learning structures. How can this be explained?

The explanation advanced here is that our experience in real world tasks is such that we develop rules and methods that seem to work reasonably well. However, these rules may be quite poor and our awareness of their inadequacy is profound. This lack of awareness exists because positive outcome feedback can occur in spite of, rather than because of, our predictive ability. As an illustration, consider the study by Wason (1960) in which he presented subjects with a three-number sequence, for example: 2, 4, 6. Subjects were required to discover the rule to which the three numbers conformed (the rule being three ascending numbers). To discover the rule, the subjects were permitted to generate sets of three numbers that the experimenter classified as conforming or not conforming to the rule. Subjects could stop at any point when they thought they had discovered the rule. The correct solution to this task should involve a search for disconfirming evidence rather than the accumulation of confirming evidence. For example, if someone believed that the rule had something to do with even numbers, this could only be tested by trying a sequence involving an odd number (i.e., accumulating vast amounts of confirming instances of even number sequences would not lead to the rule). The fact that only 6 of 29 subjects found the correct rule the first time they thought they did illustrates the dangers of induction by simple enumeration. As Wason (1960) points out, the solution to this task must involve "a willingness to attempt to falsify hypotheses, and thus to *test those intuitive ideas which so often carry the feeling of certitude* [p. 139; author's emphasis]."

It is important to emphasize that in Wason's experiment, where actions were *not* involved, a search for disconfirming evidence is possible. However, when actions are based on judgment, learning based on disconfirming evidence becomes more difficult to achieve. For example, consider how one might erroneously learn an incorrect rule for making judgments by focusing on the hypothetical case of a manager learning about his or her predictive ability concerning the potential of job candidates. The crucial factor here is that actions (e.g., accept/do not accept) are contingent on judgment. Therefore, at a subsequent date, the manager can only examine *accepted* candidates to see how many are successful. If there are many successes (which, as is shown later, is likely), these instances all confirm the rule. Indeed, the important point here is that it would be difficult to disconfirm the rule, even though it might be erroneous. One way in which the rule could be tested would be for the manager to accept a subset of those he or she judged to have low potential and then to observe their success rate. If their rate was as high as those judged to be of high potential, the rule would be disconfirmed. However, a systematic search for disconfirming evidence is rare and could be objected to on utilitarian and/or even ethical grounds; one would have to withhold the preferred action from some of those judged most deserving and give it to some judged least deserving. Therefore, utilitarian and/or ethical considerations may prevent one from even considering the collection of possible disconfirming information. Note that the tendency not to test hypotheses by disconfirming instances is a direct consequence of the task structure in which actions are taken on the basis of judgment. Furthermore, as Wason (1960) points out: "In real life there is no authority to pronounce judgment on inferences: the inferences can only be checked against the evidence [p. 139]." Therefore, large amounts of positive feedback can lead to reinforcement of a nonvalid rule.

Although outcomes contingent on the action not taken may not be sought, it is still the case that one can examine the number of positive hits and false positives as a way to check on the accuracy of one's predictions. Therefore, although such information is incomplete for accurately assessing the relationship between predictions and outcomes, such information is what most people have available. It is therefore important to consider the factors that affect these variables.

FACTORS AFFECTING POSITIVE HITS AND FALSE POSITIVES

In order to examine the number of positive hits and false positives that will result from making predictions in selection tasks, some notation is necessary. Let:

N = number of total decisions to be made; i.e., total number of "applicants."

$p(x \geqslant x_c) = \phi$ = selection ratio; i.e., the unconditional probability of receiving action A.

$p(y \geqslant y_c) = br$ = base rate; i.e., the unconditional probability of exceeding the criterion.

$p(y \geqslant y_c | x \geqslant x_c) = ph$ = positive hit rate.

$p(y < y_c | x \geqslant x_c) = fp$ = false positive rate.

ρ_{xy} = correlation between predictions and outcomes.