

The adaptive decision maker

John W. Payne

*Fuqua School of Business
Duke University*

James R. Bettman

*Fuqua School of Business
Duke University*

Eric J. Johnson

*The Wharton School
University of Pennsylvania*



CAMBRIDGE
UNIVERSITY PRESS

Published by the Press Syndicate of the University of Cambridge
The Pitt Building, Trumpington Street, Cambridge CB2 1RP
40 West 20th Street, New York, NY 10011-4211, USA
10 Stamford Road, Oakleigh, Victoria 3166, Australia

© Cambridge University Press 1993

First published 1993

Library of Congress Cataloging-in-Publication Data

Payne, John W.

The adaptive decision maker / John W. Payne, James R. Bettman,
Eric J. Johnson

p. cm.

Includes bibliographical references and indexes.

ISBN 0-521-41505-5. – ISBN 0-521-42526-3 (pbk.)

1. Decision-making. 2. Adaptability (Psychology) I. Bettman,
James R. II. Johnson, Eric J. III. Title.

BF448.P39 1993

92-21581

153.8'3-dc20

CIP

A catalog record for this book is available from the British Library.

ISBN 0-521-41505-5 hardback

ISBN 0-521-42526-3 paperback

Transferred to digital printing 2002

Contents

<i>Preface</i>	<i>page xi</i>
1 Adaptive decision behavior: An introduction	1
Flexibility in decision making	1
Factors influencing contingent decision behavior	3
The benefits and costs of contingent decision behavior	4
Overview of a framework for adaptive decision behavior	9
Research on contingent behavior in other areas of psychology	15
The organization of the book	17
2 Contingencies in decision making	20
Introduction	20
Decision strategies	22
Decision problem characteristics: Task effects	34
Decision problem characteristics: Context effects	54
Contingent usage of strategies for probabilistic reasoning	66
Summary	68
3 Deciding how to decide: An effort–accuracy framework	70
Introduction	70
Costs and benefits in decision making	72
Cognitive effort and decision strategies	75
The accuracy of decisions	88
Anticipated accuracy and effort	91
Trading accuracy for effort savings	92
Relationships to other frameworks	99
Summary	114

4 Studying contingent decisions: An integrated methodology	117
Introduction	117
Production systems as models of decision strategies	118
A Monte Carlo study of effort and accuracy in choice	123
Process-tracing methods	144
Validation by relating the simulation results to prior studies of contingent processing	147
Validation of the simulation results by direct tests	154
Summary	168
5 Constructive processes in decision making	171
Introduction	171
The nature of constructive processes	172
When will constructive processes be used?	177
Accuracy–effort tradeoffs and constructive processes	178
Noticing and exploiting regularities	180
Summary	189
6 When may adaptivity fail?	192
Introduction	192
Deficits in knowledge	193
Inability to execute strategies	214
Summary	216
7 Improving decisions and other practical matters	218
Introduction	218
Approaches to decision aiding	220
The measurement and prediction of preferences	239
Summary	247
8 The adaptive decision maker: A look backward and a look forward	248
Introduction	248
What have we learned?	248
New directions for future research	250
Summary	263
Appendix: The Mouselab system	264
Introduction	264
The mouse as pointing device	266

<i>Contents</i>	ix
Overview of the Mouselab system	268
Summary	278
<i>References</i>	279
<i>Name index</i>	309
<i>Subject index</i>	315

Adaptive decision behavior: An introduction

Human rational behavior is shaped by a scissors whose two blades are the structure of task environments and the computational capabilities of the actor.

(Simon, 1990, p. 7)

Flexibility in decision making

One of the most fascinating aspects of human decision behavior is the flexibility with which individuals respond to a wide variety of task conditions. Preference judgments, assessments of uncertainty, and choices among alternative courses of action all can be affected by minor changes in the task environment. To illustrate, imagine that you are a senior member of the faculty of a psychology department at a private university. One of your responsibilities is to help in the hiring of new faculty. One day, the chairperson of your department drops the files of two job applicants on your desk. She would like to know which one of the two job applicants you would prefer to invite in for a job interview. The files contain information on each applicant's educational background, prior publication record, current research and teaching interests, and evaluations of prior teaching performance, among other information. How would you go about processing the information about the two applicants in order to make a choice? How would you solve the choice problem if one applicant seemed to offer more potential as a teacher whereas the other applicant offered more potential as a researcher and colleague?

Now imagine that you are in the same situation as just described, except that your chairperson puts the files of a dozen applicants on your desk. She still wants you to choose the single applicant you would most prefer to bring in for a job interview. Again, some

applicants seem better on some dimensions (e.g., current research interests), whereas others seem better on other dimensions (e.g., prior teaching record). How would you go about processing the information about the dozen applicants in order to make a choice? Is your strategy for making the decision the same regardless of whether the number of alternatives is 2 or 12?

Much research suggests that your strategy for processing information will differ depending upon the number of alternatives to be considered. When faced with decision problems involving just two or three alternatives, people often use decision strategies that process all relevant information and require one to decide explicitly the extent to which one is willing to trade off less of one valued attribute or dimension (e.g., research potential) for more of another valued attribute (e.g., teaching potential). Such a decision process, involving the use of all relevant information and making explicit tradeoffs, is often associated with normative theories of preferential choice (see, e.g., Keeney & Raiffa, 1976).

When faced with more complex choice problems involving many alternatives, people often adopt simplifying (heuristic) strategies that are much more selective in the use of information. Further, the strategies adopted tend to be noncompensatory, in that excellent values on some attributes cannot compensate for poor values on other attributes. As an example, you might decide when faced with 12 applicants to eliminate from further consideration any applicant who has not had a research publication. Tversky (1972) refers to such a strategy as an elimination-by-aspects process.

The basic thesis of this book is that an individual's use of multiple decision strategies in different situations, including various simplifying methods or choice heuristics, is an adaptive response of a limited-capacity information processor to the demands of complex decision tasks. Further, we argue that the specific strategies used to solve particular decision problems are usually *intelligent* responses under the assumption that people have multiple goals for decisions, including both the desire to be accurate and the desire to conserve limited cognitive responses. Thus, we believe that how people decide how to decide is predictable when both the benefits and costs of specific decision strategies in particular task environments are taken into account and that people often select strategies that are appropriate to the circumstances.

Factors influencing contingent decision behavior

The two potential decision strategies described in the preceding example – that is, the use of explicit tradeoffs and the use of elimination rules – illustrate processing that is contingent on variations in the properties of the decision problem, such as the number of alternatives available. Throughout this book, our emphasis will be on the interaction of such properties of decision tasks and the limited processing capabilities of the decision maker in determining the strategies used and preferences we observe. This emphasis is consistent with the view expressed by Simon (1990) on the primacy of task and computational capabilities in determining human rational behavior.

However, factors other than properties of the decision task can also effect how a person decides how to solve a particular decision problem. For example, the reactions to a given problem can be moderated by a host of individual difference variables. Prior task knowledge and expertise in a problem domain represent two individual level factors, which can significantly affect how information is processed (Alba & Hutchinson, 1987; Chi, Glaser, & Farr, 1988). Shanteau (1988), for instance, argues that experts almost always follow some sort of divide-and-conquer process in which large problems are broken down into smaller parts, which are then solved, and then the partial solutions are put back together again. Individuals are also likely to differ with respect to the difficulty they experience with different types of reasoning operations (e.g., qualitative vs. quantitative). An example of a qualitative operation is determining whether an applicant has or has not had a prior research publication. An example of a more quantitative operation is determining how many additional publications would be needed to compensate for a poor prior teaching record.

Finally, decisions are generally not made in a social vacuum; rather, many social factors can influence decision making (Tetlock, 1985). For instance, even if an individual is making a decision, he or she may feel accountable to others such as family members or superiors in a business organization. Such feelings of accountability can affect how one makes a decision; Simonson (1989), for example, has shown that the need to justify a decision to others causes the choice to be more sensitive to certain aspects of the decision task.

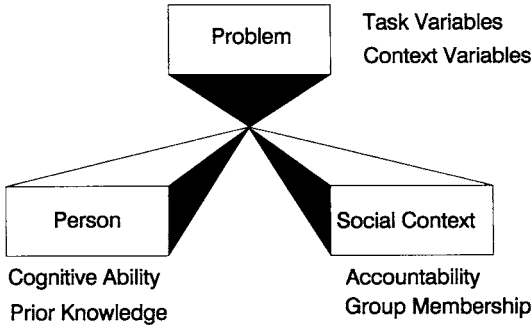


Figure 1.1. Contingent strategy selection.

Figure 1.1 illustrates the three major classes of factors that we believe influence which strategy is used to solve a particular decision problem: characteristics of the decision problem, characteristics of the person, and characteristics of the social context.

At a more detailed psychological level of analysis, these three major types of factors influencing strategy choice affect the availability, accessibility, processability, and perceived benefits of various decision strategies. For instance, prior knowledge, obtained either through experience or training, will determine which strategies are available to a decision maker in his or her memory. Experience in a decision domain also may impact the frequency and recency with which available strategies have been used, thus affecting the accessibility of various strategies. That is, experiences with strategies will affect the likelihood of recalling a particular strategy when a person faces a new decision problem. Characteristics of the problem, such as how information is displayed, can affect how much cognitive effort is needed to implement various strategies (processability). Finally, the characteristics of the social context can influence the relative importance of such factors as the justifiability of a decision in determining strategy selection.

The benefits and costs of contingent decision behavior

This brief discussion has stressed that decision behavior is contingent upon a variety of factors: Decision makers adapt to different situations. Perhaps this flexibility is not that surprising; there are potential benefits to adaptivity. At a macro level, the flexibility of

early man (e.g., willingness to eat a variety of foods) may have played a major role in the survival of the species (Calvin, 1986). At a more micro level, the flexibility of organizations can determine their chances for success in a competitive and turbulent environment (Peters, 1987). Finally, at the individual level, flexibility of response (adaptivity) is generally viewed as a mark of intelligence. Specifically, Feldman and Lindell (1990) have emphasized that flexibility of response to decision tasks is a key to the survivability of an organism. Further, they argue that "irrationality observed in any given instance is evidence of the variation in behavior that must occur if adaptation to a given environment is to take place" (pp. 107–108). Thus, a particular decision error or bias, as indicated by a deviation between behavior and a normative model, may not really be an error from a long-run, adaptive point of view. Individuals may try different behaviors and observe the results as part of a process of learning about their environment and learning how to adapt to that environment over time.

Flexibility in response may have long-run value; however, it unfortunately can also lead to short-run errors in judgment. For instance, the use of noncompensatory processes in multialternative choice can lead to the elimination of potentially good alternatives early in the decision process. Another example of how flexibility in decision making can lead to difficulties is the now classic preference reversal phenomenon (Lichtenstein & Slovic, 1971). Common sense suggests that good decisions are consistent decisions, in that small changes in the way in which a question is asked should not change what we prefer. However, Sarah Lichtenstein and Paul Slovic (1971) showed more than 20 years ago that the expressed preference order between two gambles often reverses, contingent upon whether the response requested is a direct choice between the gambles or a bidding price for each gamble. That is, the same individual would choose gamble A over gamble B and would bid more for gamble B than for gamble A, a reversal in preference. Such reversals were even replicated in a Las Vegas casino setting (Lichtenstein & Slovic, 1973), where individuals could win or lose substantial amounts of their own money.

Tversky, Sattath, and Slovic (1988) have shown more recently that people's tradeoffs between attributes (e.g., lives vs. dollars) also are contingent on the nature of the response. The more prominent dimension (i.e., lives for most people) looms larger when the

decision maker responds by making a choice as compared with when he or she responds by making a matching response, in which an aspect of one option is adjusted so that this option matches another option in overall value.¹ Hence, the tradeoff between lives and dollars is different for matching and choice. Tversky, Sattath, and Slovic suggest that choice tends to elicit qualitative types of reasoning strategies that focus on the most important attribute, whereas matching tasks elicit more quantitative types of reasoning.

If one's preferences or beliefs are affected by subtle changes in the presentation of information or changes in the way questions are asked, decision makers may be vulnerable to strategic manipulation by others. Tversky and Sattath (1979), for example, discuss how placing constraints on the order in which elements of a choice set are considered by an individual (i.e., an agenda) can affect the preference order of that individual. Thus, the flexible use of cognitive processes to make decisions, contingent on task factors, has both benefits and costs for the individual.

The highly contingent nature of decision behavior also poses problems (costs) and creates opportunities (benefits) for decision researchers. At a theoretical level, the fact that decision processes are not invariant across task environments complicates the search for a small set of underlying principles that can describe observed behavior. The research question becomes not what is *the* cognitive process used to make decisions, but instead *when* are different decision processes most likely to be used. More specifically, the cognitive control or metacognitive question of how one decides how to decide becomes crucial. This question is a major focus of our book.

The importance and pervasiveness of task and context effects also may create a view of decision research as a fragmented and chaotic field. As one answer to that problem, Hogarth and Einhorn (1992) suggest focusing on the effects of task variables on simple psychological processes like the sequential processing of information. They

¹ To illustrate a matching response, imagine that you are asked to consider the following two programs for dealing with traffic accidents, described in terms of yearly costs (in millions of dollars) and the number of casualties per year: Program X is expected to lead to 570 casualties and cost \$12 million, while program Y is expected to lead to 500 casualties and cost \$??. Your task is to provide a value for the cost of program Y, presumably some amount greater than \$12 million, that would make it equal in overall value to program X.

believe that a wide range of judgmental effects can be explained in terms of the interaction of task variables with simple information-processing strategies. We agree with Hogarth and Einhorn's emphasis on the effects of task variables on decision strategies.

The contingent nature of decision behavior has important implications at a more applied level as well. The lack of invariance across tasks that are seemingly similar (e.g., choice vs. bidding for the same lotteries) calls into question the validity of the judgmental inputs needed to operationalize such decision-aiding techniques as decision analysis (Watson & Buede, 1987). On the other hand, as noted by Tversky (1988b), it could be argued that the evidence of contingent decision processes shows that people may greatly benefit from various decision aids. He suggests that rather than abandoning decision analysis, we try to make it more responsive to the complexities and limitations of the human mind.

Because individuals adjust their decision strategies depending upon the decision task, decisions can sometimes be improved by rather straightforward, inexpensive changes to the information environments within which individuals make judgments and choices. For example, in the 1970s the provision of unit price information in supermarkets was promoted as a way of increasing consumer welfare. However, several studies showed that people were either not aware of the unit price information or were not using it. Jay Russo (1977) argued that people would like to compare alternatives directly on important attributes like unit prices; however, he also argued that it was difficult for most consumers to process unit price information because of the way in which the information was displayed. Each unit price typically was available only under each item on a shelf. Russo then argued that people would tend to ignore unit price information that was not easy to process. Thus, making information available was not sufficient to change consumer behavior; the available information also had to be processable. Russo demonstrated the power of this argument by showing that consumers' actual purchase decisions could be altered by making a simple change in the format used to present unit price information – putting all the available information on unit prices together in an easy-to-read list with unit prices ranked from lowest to highest. An important area of public policy that is currently greatly concerned with information provision issues is the communication of risk information to people about such hazards as radon levels in homes.

More generally, Arkes (1991) argues that efforts to improve (debias) intuitive judgments can be facilitated by considering the costs and benefits of various cognitive processes underlying those judgments.

Finally, the contingent nature of decision behavior has important implications for those whose job it is to measure and predict preferences. Market researchers, for example, have begun to wonder how robust their methods are to changes in decision tasks and contexts (Green & Srinivasan, 1990). More generally, recent decision research argues that preferences for objects of any complexity are often constructed – not merely revealed – in the generation of a response to a judgment or choice task (Slovic, Griffin, & Tversky, 1990). Thus, the preferences that are measured, in at least some situations, may reflect labile values, that is, values that depend on how the questions are asked.

The idea of constructive preferences means more than simply that observed choices and judgments are not the result of a reference to a master list of values in memory. Hand in hand with the notion of constructive preferences is the idea that preferences also are not necessarily generated by some consistent and task-invariant algorithm such as expected value calculation² (Tversky, Sattath, & Slovic, 1988). Instead, it appears that decision makers have a repertoire of methods for identifying their preferences. March (1978) attributes the constructiveness of preferences to limits on the information-processing capacity of individuals. In his words, “human beings have unstable, inconsistent, incompletely evoked, and imprecise goals at least in part because human abilities limit preference orderliness” (March, 1978, p. 598). A key research question is to understand what elements of the judgment or choice task cause different methods to be used in construction of the observed preferences.³

² The expected value algorithm involves the multiplication of the value (payoff) of each outcome that might occur if a particular course of action is selected by its associated probability of occurrence, adding the payoff-probability products across all the outcomes of an alternative course of action. The expected value decision rule is to select that alternative course of action with the largest expected value. As a strategy for making decisions, the expected value rule has been in existence for at least several centuries.

³ The notion of constructed preferences is consistent with the “philosophy of basic values,” which holds that people lack well-differentiated values for all but the most familiar of evaluation tasks (Fischhoff, 1991). Fischhoff draws an interesting comparison between the philosophy of basic values and the “philosophy of articulated values,” which assumes that people have values for all (most) evaluation questions and that the trick is just to ask the right question in the right way.

Overview of a framework for adaptive decision behavior

Given the extensive evidence that human decision behavior is a highly contingent form of information processing, there is a need for a framework within which such contingent behavior might be understood. A major purpose of this book is to offer such a framework. We are concerned with the task conditions and psychological mechanisms that lead to the selection of one cognitive process rather than another in solving a particular decision problem. The theoretical framework we offer allows us to answer the question of when different strategies will be used by a decision maker; we also provide evidence for that framework.

As noted earlier, our basic thesis is that the use of various decision strategies is an adaptive response of a limited-capacity information processor to the demands of complex task environments. In particular, we emphasize understanding adaptive strategy use in terms of the accuracy of and the cognitive effort required by various available strategies. That is, we believe that the two primary considerations underlying contingent decision behavior are the desire to achieve a good decision and the desire to minimize the cognitive effort needed to reach a decision. Although we believe that decision processing generally reflects reasonable effort and accuracy trade-offs, there are also important constraints on human adaptivity in decision making. People do make errors in judgments and choices. Thus two additional focuses of the book are limits on adaptivity in decision making and how decision making might be improved.

Decision strategies and problem solving

We have used the term *decision strategy* a number of times already without offering a specific definition. Within our framework, we define a decision strategy as a sequence of mental and effector (actions on the environment) operations used to transform an initial state of knowledge into a final goal state of knowledge where the decision maker views the particular decision problem as solved. For example, the cognitive operations used to transform knowledge states might include such operations as acquiring an item of information from the external environment or comparing two quantities to determine which is larger.

Included in the initial state of knowledge are facts about the

problem. For example, in the decision problem posed at the beginning of the chapter, there are a dozen job candidates available, there are uncertainties to be considered, and there is information available about how well the alternative job candidates meet various objectives, such as good teaching and good research performance. The initial state of knowledge will also include general goal statements regarding the task, such as “choose the most preferred candidate to invite for a job interview.” Of course, sometimes the initial problem state is not that well defined, and the decision maker is faced with the need to set up subgoals and evoke processes to accomplish such subtasks as the generation of new alternatives.

As the decision maker applies operators to states of knowledge, new intermediate states of knowledge about the decision problem are generated. The decision maker might learn, for instance, that the first several faculty job candidates have no prior teaching experience. As another example, after applying a set of elimination operators to the faculty hiring problem described earlier, one might reach an intermediate state of knowledge in which the original problem is transformed into a choice among only a few candidates who have prior publication records. Note that this view of decision processing argues that the problem situation is constantly being redefined by the decision maker. After the application of additional operators, an intermediate state of knowledge could be transformed into the final goal state, in which the preferred alternative has been identified.

One distinction between decision making and other types of problem-solving tasks is that decision problems are generally ill-defined about exactly how the final goal state is to be characterized. For example, at the beginning of the problem of selecting among job candidates, you may not have a good sense of how much you are prepared to trade off research potential for teaching potential, or even if such tradeoffs are required. Thus, your task of identifying which candidate best meets the goal state requires, in part, that the goal state be clarified during the decision process. Hogarth (1987) argues that people often prefer not to directly confront the conflict of trading off more of one valued attribute for less of another valued attribute, which is inherent in many decision problems. Thus, he argues that people may sometimes use noncompensatory decision strategies to solve even simple decision problems as a way to avoid conflict. Hogarth's argument points out that accuracy and effort considerations may not be the only determinants of strategy choice;

we agree. However, we believe that accuracy and effort are the *primary* determinants of contingent strategy use.

The view of decision making as the application of a series of operators to knowledge states is not unique to us. A similar conception of decision processing, for example, is offered by O. Huber (1989). More generally, our view of decision strategies is closely related to views of problem solving as the application of a sequence of mental operators (see, e.g., Newell & Simon, 1972; Holland, Holyoak, Nisbett, & Thagard, 1986).

For some decision problems, the strategy used to solve the problem will be a simple memory retrieval process. As an example, when asked the question, What is your favorite college basketball team? the answer – Duke – is drawn readily from the memories of the first two authors of this book. This type of strategy for solving decision problems is called affect referral (P. Wright, 1975). No information is processed about the characteristics of the alternatives being considered; instead, the answer is simply based on prior evaluations of the alternatives. This book is concerned, however, with the strategies people use to solve decision problems for which affect referral does not provide an acceptable solution. That is, we are interested in decision behavior when a person is faced with a decision problem of some novelty and complexity. We also assume that more than one sequence of operators (strategies) is available to the decision maker for decision problems of any complexity.

We assume further that the operators used to transform knowledge states in decision making can be represented as productions of the form IF (condition 1, . . . , condition n) THEN (action 1, . . . , action m). An example of a production might be “If there are more than four alternatives, then eliminate those alternatives that cost more than some target amount, say \$100.” The conditions can include goals and subgoals (e.g., If the goal is to please the chairperson then. . .) as well as information on problem states. The presence of goals and subgoals in the condition side of productions provides the basis for a hierarchical structure to decision behavior. Further, as acknowledged by J. Anderson (1983), the setting of goals can be used to favor special modes of processing, such as efficiency. The actions can include actions on the environment (e.g., then eliminate alternative X) and the creation of new knowledge states (e.g., then alternative Y is better on the cost attribute than alternative Z). Satisfaction of the conditions depends on the match between the conditions and

active information in a person's working memory. Like Holland et al. (1986) and others, we assume that active information may come directly from perceptual input, from the previous actions of other operators, or from a more permanent memory store.

This idea that the conditions of operators are matched against active information in working memory, coupled with the notion expressed previously that the problem situation is constantly being redefined by the decision maker, usually implies that decision making is very sequential and dynamic. That is, which operations a person performs next in solving a decision problem will depend on the information active in memory as the result of the actions of prior mental operations. Further, given that working memory has limited capacity, the information in working memory will likely reflect the most recently performed operations. As a result, the preferences we observe will depend on the particular sequence of operations used to solve a decision problem, and the performance of some operation at time t may inhibit or facilitate the performance of another operation at time $t + 1$. Kahneman and Tversky (1979b) make a related point, arguing that the order in which simplifying procedures (what they call "editing operations") are applied to a risky choice problem may permit or prevent the later application of other editing operations. As noted by Kahneman and Tversky, the sequence of operations is likely to vary as a function of variables like information display and the particular set of alternatives in the choice set. Thus, the preference order among alternatives need not be invariant across contexts.

Cognitive effort and accuracy

We also believe that the operators used in performing the types of decision tasks in which we are interested take cognitive resources to execute and that different operators may require different amounts of cognitive resources. For example, working memory limitations may often make certain kinds of cognitive processes (e.g., the mental multiplication of large numbers) very difficult to perform. There may also be circumstances in which some operators may not be feasible given the constraints of the human information-processing system. Because cognitive resources are needed to implement individual mental operations, increasing the number of operators or using more demanding operators to reach the goal state creates a more