Choices, Values, and Frames

Edited by
Daniel Kahneman
and Amos Tversky

Choices, Values, and Frames

This book presents a selection of the research that grew from the editors' early collaboration on "Prospect Theory," the landmark article that offered the first compelling alternative to the standard "rational agent" model of choice under risk. In the spirit of the highly influential volume *Judgment Under Uncertainty*, first published in 1982, this book collects numerous theoretical and empirical articles that have become classics, important extensions, and applications that range from principles of legal compensation to the behavior of New York cab drivers on busy days. Several surveys prepared especially for this volume illustrate the scope and vigor of the behavioral study of choice.

Theoretically elegant and empirically robust, the research collected in this volume represents an approach to the science of decision making that has influenced numerous fields of study, including decision theory, behavioral economics and behavioral finance, consumer psychology, the study of negotiation and conflict, medical decision making, legal analysis and practice, well-being studies, political science, and philosophical investigations of rationality and ethics. The book provides an accessible introduction to the study of decision-making behavior and is an indispensable reference source for students and specialists.

Daniel Kahneman and the late Amos Tversky have started a new perspective on the traditional economic categories of choice, decision, and value. A series of experimental and empirical studies by them and others have rejected traditional assumptions of rationality. Even more importantly, these scholars have developed alternative generalizations with significant predictive power and have found empirical verification for them. This outstanding collection of studies will make these new results readily accessible.

Kenneth J. Arrow, Nobel Memorial Laureate in Economic Science, Stanford University

---

Daniel Kahneman is Eugene Higgins Professor of Psychology and Professor of Public Affairs at the Woodrow Wilson School of Public Affairs, Princeton University.

The late Amos Tversky (1937–1996) was Davis–Brack Professor of Behavioral Science at Stanford University.

CHOICES, VALUES, AND FRAMES

Edited by
DANIEL KAHNEMAN
Princeton University

AMOS TVERSKY
Stanford University

Russell Sage Foundation

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town, Singapore, São Paulo, Delhi

Cambridge University Press
32 Avenue of the Americas, New York, NY 10013-2473, USA

www.cambridge.org
Information on this title: www.cambridge.org/9780521621724

First published 2000
10th printing 2009

Printed in the United States of America

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

Library of Congress Cataloging in Publication Data

Choices, values, and frames / edited by Daniel Kahneman, Amos Tversky
p. cm.
ISBN 0-521-62172-0 – ISBN 0-521-62749-4 (pbk.)
1. Decision making. 2. Uncertainty. 3. Risk-taking (Psychology). 4. Value.
I. Kahneman, Daniel, 1934–. / II. Tversky, Amos.
HD30.23.C469 2000
658.4'03–dc21 99-059883

ISBN 978-0-521-62172-4 hardback
ISBN 978-0-521-62749-8 paperback

Contents

Preface

This preface is not the one that Amos Tversky and I intended to write. Soon after Amos learned early in 1996 that he only had a few months to live, we decided to edit a joint book on decision making that would collect much of our work on this topic and congenial research by others. The collection was to be a sequel and companion to a volume on heuristics and biases of judgment that we had edited together with Paul Slovic many years earlier, and a substitute for a book that Amos and I had promised the Russell Sage Foundation.

Most of the editorial task was completed quickly, although some new pieces that Amos wanted to include – notably one that I was to write – were only completed long after he was gone. The problem of writing a preface was more difficult than finding articles we liked. Our initial aspirations for the preface were high; we were going to write a broad essay presenting a view of how the field had changed in the preceding 20 years. But we ran out of time before we had a presentable product. Amos advised me to "trust the model of me that is in your mind" and write for both of us. This was not advice that I was able to follow: the risk of writing in his name statements that he might have rejected proved intimidating to the point of paralysis. I chose to take the smaller risk of writing on my own and in a more personal vein than we would have adopted if we had written jointly. My intent is to provide a perspective on some dominant themes of this book by tracing them to the years that we spent working closely together on the problem of choice.

Amos and I began to collaborate in the study of decision making in 1974, soon after completing a review of our earlier work on judgment, which appeared in *Science* that year. Because I had no expertise in the formal study of choice, in which he was already a star, Amos encouraged me to read the text on mathematical psychology that he had coauthored with his teacher Clyde Coombs and his contemporary Robyn Dawes. The most relevant chapter was concerned with preferences between gambles; it included an introduction to utility theory and an exposition of the Allais paradox. Our discussions of this chapter quickly focused on a remarkable discrepancy: the theoretical analysis implied that the carriers of utility were states of wealth, but the outcomes of gambles were always described as gains and losses. Some extra assumption was clearly required to bridge the gap – either a spontaneous mental activity that transforms gains and losses into states of wealth, or a highly constrained correspondence between the psychophysics of wealth and the psychophysics of gains and losses. Neither hypothesis seemed appealing. This observation eventually led us to develop a theory of choice under risk that we called prospect theory (Chapter 2).

Our method of research in those early Jerusalem days was pure fun. We would meet every afternoon for several hours, which we spent inventing interesting pairs of gambles and observing our own intuitive preferences. If we agreed on the same choice, we provisionally assumed that it was characteristic of humankind and went on to investigate its theoretical implications, leaving serious verification for later. This unusual mode of empirical research enabled us to move quickly. In a few giddy months we raced through more than twenty diverse theoretical formulations. Among others, we considered a treatment of risky choice in terms of regret, but we eventually abandoned this approach because it did not elegantly accommodate the pattern of results that we labeled 'reflection': changing the signs of all outcomes in a gamble almost invariably changed the direction of preferences from risk averse to risk seeking or vice versa.

Prospect theory was nearly complete in the spring of 1975; it was then called value theory. A reader familiar with the final version would find almost all the important ideas in that early draft. However, the three additional years that we spent fine-tuning the theory turned out to be essential to its viability. The idea of editing operations, for example, was our answer to many potential counterexamples; it took three years to get to it. Amos was capable of infinite patience. His frequently repeated exhortation "let's get it right" and the contagious pleasure that he found in the activity of thinking enabled us to persevere in the task of anticipating possible objections – even minor ones – long enough to ensure that they had been resolved, not papered over.

The theory that we constructed was as conservative as possible. We stayed within the decision theoretic framework in which choice between gambles is the model for all decisions. We did not challenge the philosophical analysis of choices in terms of beliefs and desires that underlies utility theory, nor did we question the normative models of rational choice offered by von Neumann and Morgenstern and later by Savage. The goal we set for ourselves was to assemble the minimal set of modifications of expected utility theory that would provide a descriptive account of everything we knew about a severely restricted class of decisions: choices between simple monetary gambles with objectively specified probabilities and at most two nonzero outcomes. Without additional assumptions, prospect theory is not applicable to gambles that have a larger number of outcomes, to gambles on events, or to transactions other than choice; it does not even specify a selling price for monetary gambles. And our hopeful statement that "the extension of equations (1) and (2) to prospects with any number of outcomes is straightforward" (p. 74) turned out to be very optimistic. We got to it some thirteen years later (see Chapter 3), and the extension was not at all straightforward.

It is fair to ask, What is the point of investing so much effort in a theory if its domain of application is so restricted and artificial? The answer is that choice between gambles is the fruit fly of decision theory. It is a very simple case, which contains many essential elements of much larger problems. As with the fruit fly, we study gambles in the hope that the principles that govern the simple case will extend in recognizable form to complex situations. A theory of choice should therefore be evaluated by two distinct criteria, and a successful theory must satisfy both. One requirement is to "get it right" within a precisely specified domain. The theory should be refutable in that domain, and it should be unrefuted. The larger prize, however, lies elsewhere. The principles of the theory should provide a heuristic benefit in the analysis of more complex decisions, by suggesting hypotheses and by providing templates and labels for the identification of phenomena. For example, although it is surely futile to "test" prospect theory against utility theory in the domain of international relations, the concepts of loss aversion and pseudocertainty are useful tools for understanding strategic decisions. No warranty is implied, of course. The scholars who use the tools to explain more complex decisions do so at their own risk.

Many chapters in the present collection explore one or another of four themes that emerged from our attempt to construct a viable theory of our chosen domain. When we published prospect theory we had a clear view of only the first two of these ideas: the nonlinearity of decision weights and the reference-dependent characteristics of the value function. The third theme, the significance of framing effects, was present in rudimentary form in that article and was articulated soon afterward. The fourth idea, the need to distinguish experienced utility from decision utility, came up early in our conversations but was only developed much later.

NONLINEAR DECISION WEIGHTS

The expectation principle of utility theory requires a linear response to variations of probability. As a descriptive generalization, this idea is transparently wrong: intuition suggests, and experiments readily confirm, that raising the probability of an outcome from .39 to .40 has much less impact on preferences than increasing the probability of the same outcome from 0 to .01 or from .99 to 1.00. This observation, which Ward Edwards and others had made before us, was incorporated in the decision weight function of prospect theory. Several chapters of this book document the progress that was subsequently achieved in the understanding of decision weights. These include the rank-dependent formulation that generalizes prospect theory to gambles with any number of outcomes (Chapter 3), Dražen Prelec's novel analysis of the shape of decision weights (Chapter 4), and an important extension to the case of gambles on events, which sheds light on both the Ellsberg and the Allais paradoxes (Chapter 5). I believe that in his last year Amos was close to achieving a major generalization of prospect theory in a unified treatment of risk and uncertainty

(Chapter 6). With characteristic restraint, however, he rejected the suggestion to label this work "extended prospect theory." Chapter 6 also illustrates an important observation that Amos and his students had analyzed in the context of support theory: the large influence of the "packing" or "unpacking" of events on subjective probabilities, and through them, on decision weights. To maintain unity of style Amos chose not to include in the present collection the results of the mathematical explorations of prospect theory to which he devoted a significant effort in the 1990s in a fruitful collaboration with Peter Wakker.

REFERENCE-DEPENDENCE AND LOSS AVERSION

Standard applications of utility theory assume that the outcomes of risky prospects are evaluated as states of wealth. This assumption was the cornerstone of the version of utility theory that Daniel Bernoulli offered in 1738, and it has been retained ever since. The proposition that the carriers of utility are states of wealth is accepted as a matter of course in economic analyses and in the prescriptions of decision analysts. However, casual observation suggests that this assumption must also be modified. In the vernacular of decision making, financial outcomes are almost always described as gains and losses; states of wealth are rarely mentioned unless death or ruin is a possibility. The argument appears to have been closed by Matthew Rabin's demonstration that no utility function for wealth can accommodate the extreme risk aversion that people exhibit when they face gambles with small stakes (Chapter 11).

The isolation effect described in prospect theory was intended as a direct experimental test of the hypothesis that preferences for gambles are determined by the utility of wealth. Respondents were asked to imagine that they had been given an unconditional gift of cash and that they now faced a choice between a sure outcome and a fair gamble. There were two versions of the problem with different combinations of the size of the gift and of the possible outcomes of the choice. The options described in the two versions were strictly identical in terms of possible states of final wealth and their probabilities, but one of the choices was between a sure gain and a positive gamble; the other was between a sure loss and a negative gamble. As expected from the reflection effect, most respondents were risk averse in the first version, and most were risk seeking in the second – much as they would have been if the initial cash gift had not been mentioned at all. This demonstration simultaneously provided a counterexample to the idea that the carriers of utility are states of wealth and positive evidence for the role of gains and losses. It was also the first application of the method that we used later to study framing effects.

The idea that the effective carriers of utility are gains and losses was not new. Markowitz had proposed it much earlier. The distinctive contribution of prospect theory was the S-shaped value function that is reproduced in many articles in this collection. The convex–concave shape contributed to the explanation of the reflection effect, which we then considered our most remarkable experimental result. The value function is also much steeper in the domain of losses than in the domain of gains, a characteristic that we later labeled loss aversion. The assumption of loss aversion helped explain two salient facts of risky choice: the almost universal rejection of gambles that offer equal probabilities to win and lose the same amount, and the increase of this aversion with the size of the stakes. We introduced loss aversion reluctantly. As we well knew, it is hardly elegant for a theory of choice to invoke different explanations of risk aversion for positive prospects and for prospects in which losses are possible, but we found no simpler model that would account for the facts. We realized only much later that loss aversion is the element of prospect theory that has the richest implications beyond its narrow domain.

We had considerable help from a friend in appreciating the broader significance of loss aversion. While we were still working on the final version of prospect theory, a young economist named Richard Thaler sought us out. Dick had begun a career of ironic and deep comments on his discipline while still a graduate student, collecting amusing anecdotal examples of everyday behaviors of consumers that violated basic assumptions of standard economic theory. He realized (Chapter 15) that many of these anomalies could be explained by extending the idea of loss aversion to riskless decisions. The most important application was to the endowment effect, which Thaler illustrated by the memorable example of a bottle of old wine which the owner would refuse to sell for $200 but would not pay as much as $100 to replace. This pattern is odd in the context of standard economic theory in which an individual's buying price and selling price for a good are assumed to differ only because of transaction costs and an income effect. In contrast, the endowment effect is readily explained by two assumptions derived from prospect theory. First, the carriers of utility are not states (owning or not owning the wine), but changes: getting the wine or giving it up. And giving up is weighted more than getting, by loss aversion.

The idea that "losses loom larger than gains" is a major theme of this collection. Loss aversion is invoked to explain phenomena as diverse as indifference curves that cross (Chapter 9), principles of legal compensation (Chapter 24), rules of commercial fairness (Chapter 18), the equity premium puzzle in financial markets (Chapter 17), and the number of hours that New York cab drivers choose to work on busy days (Chapter 20), among many others. There is no reason to suspect that the topic is exhausted. The ramifications of the hypothesis that preferences are reference dependent and loss averse are yet to be fully explored.

FRAMING EFFECTS AND MENTAL ACCOUNTING

Another theme of this collection is the dependence of choices on the description and interpretation of decision problems. Amos and I turned to the study of framing immediately after the completion of prospect theory. Most of our evidence was at hand in 1979 and was first published in 1981. Later essays (Chapters 1 and 3) revisited the same results, reflecting our evolving understanding of their significance as a challenge to the general model of rational choice.

A significant and perhaps unfortunate early decision concerned the naming of the new concept. For reasons of conceptual and terminological economy we chose to apply the label "frame" to descriptions of decision problems at two levels: the formulation to which decision makers are exposed is called a frame and so is the interpretation that they construct for themselves. Thus, framing is a common label for two very different things: an experimental manipulation and a constituent activity of decision making. Our terminological parsimony was helpful in securing the acceptance of the concept of framing, but it also had its costs. The use of a single term blurred the important distinction between what decision makers do and what is done to them: the activities of editing and mental accounting on the one hand and the susceptibility to framing effects on the other.

Prospect theory includes a set of rules of editing for simple gambles. We assumed a preliminary phase of decision making in which specified editing operations transform the problem – usually into a simpler form. The initial motivation for introducing editing operations was defensive: they eliminated some foolish predictions to which prospect theory seemed otherwise committed. For example, the properties of decision weights would imply a preference for the prospect ($100, .01; $100, .01) over ($100, .02). The prediction is wrong because most decision makers will spontaneously edit the former prospect into the latter and treat them as equivalent in subsequent operations of evaluation and choice. Editing operations provided an explicit and psychologically plausible defense against such superficial counterexamples to the core of the theory. Although we did not immediately see it, the conception of editing also led to the more general observation that the true objects of evaluation and choice are neither objects in the real world nor verbal descriptions; they are mental representations. This conceptual move was novel in the context of decision research, but it is entirely natural for cognitive psychologists. Anyone who has taken a course in perception has learned to distinguish objective reality from the proximal stimulus to which the observer is exposed and to distinguish both reality and the stimulus from the mental representation that the observer eventually constructs.

Richard Thaler's early ideas about mental accounting (Chapter 15), which we adopted and elaborated (Chapter 1) offered an informal treatment of how people organize decisions and outcomes by lumping some together and segregating others, in ways that often violate standard assumptions of economic theory. Much of what is known about the psychology of active framing is discussed by Thaler in Chapter 14. A particularly significant feature of the accounting metaphor is that mental accounts are eventually closed and that strong emotions may be experienced at those times of reckoning. An implication of this insight is that people can to some extent control their own rewards and punishments by choosing whether to close an account or keep it open as well as deciding when to evaluate it. Variations on this powerful theme help explain why sunk costs are not ignored (Chapters 14 and 15), why cab drivers stop work too early on busy days (Chapter 20), why investors are more likely to sell "winners" than "losers" when they lighten their portfolio of stocks (Chapter 21), and why financial returns are much higher on stocks than on bonds (Chapter 17).

As is evident in Chapter 1, we were initially more interested in framing effects than in the activity of framing. We were surprised by how easy it was to construct different versions of a decision problem that were transparently equivalent when considered together but evoked different preferences when considered separately. Framing effects are inherently interesting, but the psychological analysis of these effects is awkward because the object of explanation is something that decision makers do *not* do: they do not spontaneously generate a common representation for decision problems that they would judge to be equivalent. Why is this so? The unexciting answer is that decision makers are generally quite passive and therefore inclined to accept any frame to which they are exposed. Framing effects are less significant for their contribution to psychology than for their importance in the real world (a recurrent theme in this book, e.g., in Chapters 13, 16, 19, and 25) and for the challenge they raise to the foundations of a rational model of decision making (Chapter 12).

The isolation effect was our first application of a method in which alternative versions of "the same" decision problem evoked different preferences. We constructed a pair of problems that were identical in the context of a particular interpretation of utility theory and suggested that the observation of different preferences was a counterexample to that theory. However, we eventually adopted a less theory-bound view of what makes two problems the same. It is the decision maker who should determine, after due consideration of both problems, whether the differences between them are sufficiently consequential to justify different choices. Violations of this lenient form of invariance demonstrate incoherence without a need for any judgment from on high about what is truly equivalent. The ubiquity of framing effects demonstrates that the human mind is not designed to achieve coherence. It took us several years to realize that violations of invariance challenge the excessively demanding conception of rationality that prevails in economics and in other sciences of decision (Chapter 12). Violations of invariance provide a compelling reason to separate descriptive from normative models of choice. It is surely rational to treat identical problems identically, but often people do not.

This was as far as Amos and I went together, but over the next decade we each explored aspects of these issues further with various other collaborators. Amos remained fascinated by the idea that the objects of evaluation are descriptions, not their referents. The notion that the carriers of subjective probability are descriptions of events was used to powerful effect in support theory, the model of probabilistic judgment on which he was still working at the end of his life (see Chapter 6). More directly, Amos expanded the concept of violations of invariance from framing effects to diverse types of preference reversals (Chapter 28). In various collaborations with close friends and colleagues, notably Paul Slovic, Eldar Shafir, Itamar Simonson, and Shmuel Sattath (Chapters 28, 29, and 34), he explored other situations in which people's preferences are influenced by variables that they would wish to ignore. My own interest in the area was in framing as an activity, and especially in the notable proclivity of decision makers to frame problems of judgment and choice more narrowly than they should (Chapter 22).

The chapters collected here only provide a glimpse at the evidence for the extreme sensitivity of choices to formulation, context, and procedure. A growing body of findings supports a radical challenge to the assumption, central to much economic theory, that stable preferences exist. The image of a decision maker who makes choices by consulting a preexisting preference order appears increasingly implausible. The alternative image is of a decision maker who chooses reluctantly and with difficulty (Chapter 34) and who constructs preferences in the context and in the format required by a particular situation (Chapters 27–29). Of course, no one wishes to pursue the idea of context dependence to the point of nihilism. Choices are not nearly as coherent as the notion of a preference order would suggest, but they are also far from random. Some explanation of their limited coherence is therefore required. Perhaps people are better described as having attitudes than as having preferences (Chapters 35 and 36).

EXPERIENCED UTILITY

Sometime in the early years of our work on the shape of the value function of prospect theory, I posed to Amos two puzzles that still fascinate me. Consider an individual who faces a course of daily injections that he expects to remain equally painful from day to day. The first injection is due tomorrow. Will the individual be willing to pay the same amount today to eliminate the last injection in a series of 20 daily injections or in a series of 5 injections? And what if a decision must be made today about a series of 5 injections that starts in 15 days? Our intuitive answers to these questions were both definite and clearly unreasonable. They were also widely shared, as we eventually learned from research by Herrnstein, Thaler, Loewenstein, and Prelec. Most people, of course, will pay more to reduce the length of a series of injections from 5 to 4 than they will pay, *ex ante*, to reduce it from 20 to 19. They will also show different patterns of declining willingness to pay to reduce a series of injections that begin tomorrow or in fifteen days. Some investigators, represented here by Loewenstein and Prelec (Chapters 32 and 33), have focused their attention on the remarkable contrast between observed time preferences and the standard economic analyses of discounting. The question that intrigued us was different: What is the justification of *any* departure from linearity in the patient's valuation of injections if the patient believes that the injections will be equally painful from day to day? More generally, what is the normative standing of sharply curved utility functions and of steep discount functions in such cases?

The injection story contains the seeds of several ideas that were only taken up much later and are yet to be fully developed. First, it assumed the possibility of eliciting the "predicted utility" of consequences independently of their "decision utility" by obtaining, in this case, the patient's beliefs about the pain of injections. The task of predicting future tastes and the relation between predicted utility and decision utility are investigated in two chapters in this collection (Chapters 40 and 41). Second, the direct measurement of the patient's pain, his or her "experienced utility," is a potential criterion for evaluating decisions. By this substantive criterion, the standard assumption that people maximize utility is not tautological but false (Chapter 42). Third, considerations of the utility of outcomes as they will actually be experienced highlights an important deficiency in the standard decision-theoretic definition of rationality. The coherence view is not only unreasonably demanding, as was noted in the earlier discussion of framing. When an objective observer considers the consequences actually experienced, the coherence criterion also appears much too permissive. The rule that any utility function is acceptable if only it is applied consistently allows too many foolish decisions to be considered rational. The chapters in the last section of the book indicate some of the progress that has been made in pursuing these ideas.

The concluding section of Chapter 1 briefly mentions the distinction between decision utility and experienced utility and hints that the latter should be the criterion for the former. Amos and I never worked together on this topic, though we often discussed it. Our views sometimes differed, especially when the discussion touched on the role of memory in life, which Amos considered more important than I did – perhaps because his memory was so much better than mine. In his chapter with Dale Griffin (Chapter 39), Amos described life as the gradual accumulation of an endowment of memories, which is not at all the view I take (Chapter 37). One of my enduring regrets is that we never resolved the difference by studying it together. We would have come closer to "doing it right," and it would have been a joy.

Daniel Kahneman
Princeton
August 1999

List of Contributors

Daniel Alder, Department of Social and Decision Sciences, Carnegie Mellon University

Linda Babcock, The Heinz School, Carnegie Mellon University

Ian Bateman, Centre for Social and Economic Research, Department of Environmental Sciences, University of East Anglia, UK

Shlomo Benartzi, John Anderson School of Management, Department of Accounting, University of California, Los Angeles

Colin F. Camerer, Division of Social Sciences, California Institute of Technology

David Cohen, Psychology Department, University of Texas, Austin

Peter Diamond, Department of Economics, Massachusetts Institute of Technology

Baruch Fischoff, Department of Social and Decision Sciences, Carnegie Mellon University

Craig R. Fox, Fuqua School of Business, Duke University

Dale Griffin, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK

John Hershey, Department of Operations and Information Management, The Wharton School, University of Pennsylvania

Christopher K. Hsee, Graduate School of Business, University of Chicago

Eric J. Johnson, School of Business, Columbia University

Daniel Kahneman, Woodrow Wilson School of Public Affairs and Department of Psychology, Princeton University

Jack L. Knetsch, Department of Economics and Resources Management, Simon Fraser University, Canada

Howard Kunreuther, Department of Operations and Information Management, The Wharton School, University of Pennsylvania

Dan Lovallo, Australian Graduate School of Management, University of New South Wales, Australia

George Loewenstein, Department of Social and Decision Sciences, Carnegie Mellon University

Jacqueline Meszaros, Business Administration Program, University of Washington, Bothell

Alistair Munro, Department of Economic and Social Studies, University of East Anglia, UK

Terrance Odean, Graduate School of Management, University of California, Davis

Dražen Prelec, The Sloan School and Department of Psychology, Massachusetts Institute of Technology

George A. Quattrone, Department of Psychology, Stanford University

Matthew Rabin, Department of Economics, University of California, Berkeley

Bruce Rhodes, Department of Economic and Social Studies, University of East Anglia, UK

Ilana Ritov, School of Education, Hebrew University, Israel

Shmuel Sattath, Hebrew University, Israel

David Schkade, Department of Management, University of Texas, Austin

Eldar Shafir, Department of Psychology, Princeton University

Itamar Simonson, Graduate School of Business, Stanford University

Chris Starmer, School of Economic and Social Studies, University of East Anglia, UK

Robert Sugden, Department of Economic and Social Studies, University of East Anglia, UK

Richard H. Thaler, Graduate School of Business, University of Chicago

Amos Tversky (Deceased), Department of Psychology, Stanford University

1. Choices, Values, and Frames

Daniel Kahneman and Amos Tversky

ABSTRACT. We discuss the cognitive and the psychophysical determinants of choice in risky and riskless contexts. The psychophysics of value induce risk aversion in the domain of gains and risk seeking in the domain of losses. The psychophysics of chance induce overweighting of sure things and of improbable events, relative to events of moderate probability. Decision problems can be described or framed in multiple ways that give rise to different preferences, contrary to the invariance criterion of rational choice. The process of mental accounting, in which people organize the outcomes of transactions, explains some anomalies of consumer behavior. In particular, the acceptability of an option can depend on whether a negative outcome is evaluated as a cost or as an uncompensated loss. The relation between decision values and experience values is discussed.

Making decisions is like speaking prose – people do it all the time, knowingly or unknowingly. It is hardly surprising, then, that the topic of decision making is shared by many disciplines, from mathematics and statistics, through economics and political science, to sociology and psychology. The study of decisions addresses both normative and descriptive questions. The normative analysis is concerned with the nature of rationality and the logic of decision making. The descriptive analysis, in contrast, is concerned with people's beliefs and preferences as they are, not as they should be. The tension between normative and descriptive considerations characterizes much of the study of judgment and choice.

Analyses of decision making commonly distinguish risky and riskless choices. The paradigmatic example of decision under risk is the acceptability of a gamble that yields monetary outcomes with specified probabilities. A typical riskless decision concerns the acceptability of a transaction in which a good or a service is exchanged for money or labor. In the first part of this article

This article was originally presented as a Distinguished Scientific Contributions Award address at the meeting of the American Psychological Association, Anaheim, California, August 1983. This work was supported by grant NR 197-058 from the U.S. Office of Naval Research.

Originally published in *American Psychologist*, 39:4, 341–50.

we present an analysis of the cognitive and psychophysical factors that determine the value of risky prospects. In the second part we extend this analysis to transactions and trades.

RISKY CHOICE

Risky choices, such as whether or not to take an umbrella and whether or not to go to war, are made without advance knowledge of their consequences. Because the consequences of such actions depend on uncertain events such as the weather or the opponent's resolve, the choice of an act may be construed as the acceptance of a gamble that can yield various outcomes with different probabilities. It is therefore natural that the study of decision making under risk has focused on choices between simple gambles with monetary outcomes and specified probabilities in the hope that these simple problems will reveal basic attitudes toward risk and value.

We shall sketch an approach to risky choice that derives many of its hypotheses from a psychophysical analysis of responses to money and to probability. The psychophysical approach to decision making can be traced to a remarkable essay that Daniel Bernoulli published in 1738 (Bernoulli 1738/1954) in which he attempted to explain why people are generally averse to risk and why risk aversion decreases with increasing wealth. To illustrate risk aversion and Bernoulli's analysis, consider the choice between a prospect that offers an 85% chance to win $1000 (with a 15% chance to win nothing) and the alternative of receiving $800 for sure. A large majority of people prefer the sure thing over the gamble, although the gamble has higher (mathematical) expectation. The expectation of a monetary gamble is a weighted average, where each possible outcome is weighted by its probability of occurrence. The expectation of the gamble in this example is $.85 \times \$1000 + .15 \times \$0 = \$850$, which exceeds the expectation of $800 associated with the sure thing. The preference for the sure gain is an instance of risk aversion. In general, a preference for a sure outcome over a gamble that has higher or equal expectation is called risk aversion, and the rejection of a sure thing in favor of a gamble of lower or equal expectation is called risk seeking.

Bernoulli suggested that people do not evaluate prospects by the expectation of their monetary outcomes, but rather by the expectation of the subjective value of these outcomes. The subjective value of a gamble is again a weighted average, but now it is the subjective value of each outcome that is weighted by its probability. To explain risk aversion within this framework, Bernoulli proposed that subjective value, or utility, is a concave function of money. In such a function, the difference between the utilities of $200 and $100, for example, is greater than the utility difference between $1,200 and $1,100. It follows from concavity that the subjective value attached to a gain of $800 is more than 80% of the value of a gain of $1,000. Consequently, the concavity of the utility function entails a risk averse preference for a sure gain of $800 over an 80% chance to win $1,000, although the two prospects have the same monetary expectation.

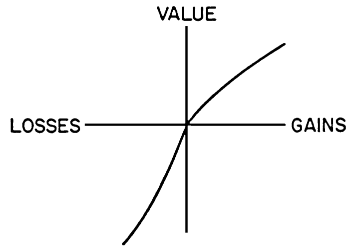

Figure 1.1. A hypothetical value function.

It is customary in decision analysis to describe the outcomes of decisions in terms of total wealth. For example, an offer to bet $20 on the toss of a fair coin is represented as a choice between an individual's current wealth W and an even chance to move to W + $20 or to W − $20. This representation appears psychologically unrealistic: People do not normally think of relatively small outcomes in terms of states of wealth but rather in terms of gains, losses, and neutral outcomes (such as the maintenance of the status quo). If the effective carriers of subjective value are changes of wealth rather than ultimate states of wealth, as we propose, the psychophysical analysis of outcomes should be applied to gains and losses rather than to total assets. This assumption plays a central role in a treatment of risky choice that we called prospect theory (Kahneman & Tversky, 1979). Introspection as well as psychophysical measurements suggest that subjective value is a concave function of the size of a gain. The same generalization applies to losses as well. The difference in subjective value between a loss of $200 and a loss of $100 appears greater than the difference in subjective value between a loss of $1,200 and a loss of $1,100. When the value functions for gains and for losses are pieced together, we obtain an S-shaped function of the type displayed in Figure 1.1.

The value function shown in Figure 1.1 is (a) defined on gains and losses rather than on total wealth, (b) concave in the domain of gains and convex in the domain of losses, and (c) considerably steeper for losses than for gains. The last property, which we label *loss aversion*, expresses the intuition that a loss of $X is more aversive than a gain of $X is attractive. Loss aversion explains people's reluctance to bet on a fair coin for equal stakes: The attractiveness of the possible gain is not nearly sufficient to compensate for the aversiveness of the possible loss. For example, most respondents in a sample of undergraduates refused to stake $10 on the toss of a coin if they stood to win less than $30.

The assumption of risk aversion has played a central role in economic theory. However, just as the concavity of the value of gains entails risk aversion, the convexity of the value of losses entails risk seeking. Indeed, risk seeking in losses is a robust effect, particularly when the probabilities of loss are substantial. Consider, for example, a situation in which an individual is forced to choose between an 85% chance to lose $1,000 (with a 15% chance to lose nothing) and a sure loss of $800. A large majority of people express a preference for the gamble over the sure loss. This is a risk seeking choice because the expectation of the gamble (−$850) is inferior to the expectation of the sure loss (−$800). Risk seeking in the domain of losses has been confirmed by several investigators (Fishburn & Kochenberger, 1979; Hershey & Schoemaker, 1980; Payne, Laughhunn, & Crum, 1980; Slovic, Fischhoff, & Lichtenstein, 1982).

It has also been observed with nonmonetary outcomes, such as hours of pain (Eraker & Sox, 1981) and loss of human lives (Fischhoff, 1983; Tversky, 1977; Tversky & Kahneman, 1981). Is it wrong to be risk averse in the domain of gains and risk seeking in the domain of losses? These preferences conform to compelling intuitions about the subjective value of gains and losses, and the presumption is that people should be entitled to their own values. However, we shall see that an S-shaped value function has implications that are normatively unacceptable.

To address the normative issue we turn from psychology to decision theory. Modern decision theory can be said to begin with the pioneering work of von Neumann and Morgenstern (1947), who laid down several qualitative principles, or axioms, that should govern the preferences of a rational decision maker. Their axioms included transitivity (if A is preferred to B and B is preferred to C, then A is preferred to C), and substitution (if A is preferred to B, then an even chance to get A or C is preferred to an even chance to get B or C), along with other conditions of a more technical nature. The normative and the descriptive status of the axioms of rational choice have been the subject of extensive discussions. In particular, there is convincing evidence that people do not always obey the substitution axiom, and considerable disagreement exists about the normative merit of this axiom (e.g., Allais & Hagen, 1979). However, all analyses of rational choice incorporate two principles: *dominance* and *invariance*. Dominance demands that if prospect A is at least as good as prospect B in every respect and better than B in at least one respect, then A should be preferred to B. Invariance requires that the preference order between prospects should not depend on the manner in which they are described. In particular, two versions of a choice problem that are recognized to be equivalent when shown together should elicit the same preference even when shown separately. We now show that the requirement of invariance, however elementary and innocuous it may seem, cannot generally be satisfied.

Framing of Outcomes

Risky prospects are characterized by their possible outcomes and by the probabilities of these outcomes. The same option, however, can be framed or described in different ways (Tversky & Kahneman, 1981). For example, the possible outcomes of a gamble can be framed either as gains and losses relative to the status quo or as asset positions that incorporate initial wealth. Invariance requires that such changes in the description of outcomes should not alter the preference order. The following pair of problems illustrates a violation of this requirement. The total number of respondents in each problem is denoted by N, and the percentage who chose each option is indicated in parentheses.

Problem 1 ($N = 152$): Imagine that the U.S. is preparing for the outbreak of an unusual Asian disease, which is expected to kill 600 people. Two alternative programs to combat the disease have been proposed. Assume that the exact scientific estimates of the consequences of the programs are as follows:

If Program A is adopted, 200 people will be saved. (72%)

If Program B is adopted, there is a one-third probability that 600 people will be saved and a two-thirds probability that no people will be saved. (28%)

Which of the two programs would you favor?

The formulation of Problem 1 implicitly adopts as a reference point a state of affairs in which the disease is allowed to take its toll of 600 lives. The outcomes of the programs include the reference state and two possible gains measured by the number of lives saved. As expected, preferences are risk averse: A clear majority of respondents prefer saving 200 lives for sure over a gamble that offers a one-third chance of saving 600 lives. Now consider another problem in which the same cover story is followed by a different description of the prospects associated with the two programs:

Problem 2 ($N = 155$): If Program C is adopted, 400 people will die. (22%)

If Program D is adopted, there is a one-third probability that nobody will die and a two-thirds probability that 600 people will die. (78%)

It is easy to verify that options C and D in Problem 2 are undistinguishable in real terms from options A and B in Problem 1, respectively. The second version, however, assumes a reference state in which no one dies of the disease. The best outcome is the maintenance of this state and the alternatives are losses measured by the number of people that will die of the disease. People who evaluate options in these terms are expected to show a risk seeking preference for the gamble (option D) over the sure loss of 400 lives. Indeed, there is more risk seeking in the second version of the problem than there is risk aversion in the first.

The failure of invariance is both pervasive and robust. It is as common among sophisticated respondents as among naive ones, and it is not eliminated even when the same respondents answer both questions within a few minutes. Respondents confronted with their conflicting answers are typically puzzled. Even after rereading the problems, they still wish to be risk averse in the "lives saved" version; they wish to be risk seeking in the "lives lost" version; and they also wish to obey invariance and give consistent answers in the two versions. In their stubborn appeal, framing effects resemble perceptual illusions more than computational errors.

The following pair of problems elicits preferences that violate the dominance requirement of rational choice.

Problem 3 ($N = 86$): Choose between:

E. 25% chance to win $240 and 75% chance to lose $760 (0%)
F. 25% chance to win $250 and 75% chance to lose $750 (100%)

It is easy to see that F dominates E. Indeed, all respondents chose accordingly.

Problem 4 ($N = 150$): Imagine that you face the following pair of concurrent decisions. First examine both decisions, then indicate the options you prefer.

Decision (i) Choose between:
A. a sure gain of $240 (84%)
B. 25% chance to gain $1000 and 75% chance to gain nothing (16%)
Decision (ii) Choose between:
C. a sure loss of $750 (13%)
D. 75% chance to lose $1000 and 25% chance to lose nothing (87%)

As expected from the previous analysis, a large majority of subjects made a risk averse choice for the sure gain over the positive gamble in the first decision, and an even larger majority of subjects made a risk seeking choice for the gamble over the sure loss in the second decision. In fact, 73% of the respondents chose A and D and only 3% chose B and C. The same pattern of results was observed in a modified version of the problem, with reduced stakes, in which undergraduates selected gambles that they would actually play.

Because the subjects considered the two decisions in Problem 4 simultaneously, they expressed in effect a preference for A and D over B and C. The preferred conjunction, however, is actually dominated by the rejected one. Adding the sure gain of $240 (option A) to option D yields 25% chance to win $240 and 75% to lose $760. This is precisely option E in Problem 3. Similarly, adding the sure loss of $750 (option C) to option B yields a 25% chance to win $250 and 75% chance to lose $750. This is precisely option F in Problem 3. Thus, the susceptibility to framing and the S-shaped value function produce a violation of dominance in a set of concurrent decisions.

The moral of these results is disturbing: Invariance is normatively essential, intuitively compelling, and psychologically unfeasible. Indeed, we conceive only two ways of guaranteeing invariance. The first is to adopt a procedure that will transform equivalent versions of any problem into the same canonical representation. This is the rationale for the standard admonition to students of business, that they should consider each decision problem in terms of total assets rather than in terms of gains or losses (Schlaifer, 1959). Such a representation would avoid the violations of invariance illustrated in the previous problems, but the advice is easier to give than to follow. Except in the context of possible ruin, it is more natural to consider financial outcomes as gains and losses rather than as states of wealth. Furthermore, a canonical representation of risky prospects requires a compounding of all outcomes of concurrent decisions (e.g., Problem 4) that exceeds the capabilities of intuitive computation even in simple problems. Achieving a canonical representation is even more difficult in other contexts such as safety, health, or quality of life. Should we advise people to evaluate the consequence of a public health policy (e.g., Problems 1 and 2) in terms of overall mortality, mortality due to diseases, or the number of deaths associated with the particular disease under study?

Another approach that could guarantee invariance is the evaluation of options in terms of their actuarial rather than their psychological consequences. The actuarial criterion has some appeal in the context of human lives, but it is clearly inadequate for financial choices, as has been generally recognized at least since Bernoulli, and it is entirely inapplicable to outcomes that lack an objective metric. We conclude that frame invariance cannot be expected to hold and that a sense of confidence in a particular choice does not ensure that the same choice would be made in another frame. It is therefore good practice to test the robustness of preferences by deliberate attempts to frame a decision problem in more than one way (Fischhoff, Slovic, & Lichtenstein, 1980).

The Psychophysics of Chances

Our discussion so far has assumed a Bernoullian expectation rule according to which the value, or utility, of an uncertain prospect is obtained by adding the utilities of the possible outcomes, each weighted by its probability. To examine this assumption, let us again consult psychophysical intuitions. Setting the value of the status quo at zero, imagine a cash gift, say of $300, and assign it a value of one. Now imagine that you are only given a ticket to a lottery that has a single prize of $300. How does the value of the ticket vary as a function of the probability of winning the prize? Barring utility for gambling, the value of such a prospect must vary between zero (when the chance of winning is nil) and one (when winning $300 is a certainty).

Intuition suggests that the value of the ticket is not a linear function of the probability of winning, as entailed by the expectation rule. In particular, an increase from 0% to 5% appears to have a larger effect than an increase from 30% to 35%, which also appears smaller than an increase from 95% to 100%. These considerations suggest a category-boundary effect: A change from impossibility to possibility or from possibility to certainty has a bigger impact than a comparable change in the middle of the scale. This hypothesis is incorporated into the curve displayed in Figure 1.2, which plots the weight attached to an event as a function of its stated numerical probability. The most salient feature of Figure 1.2 is that decision weights are regressive with respect to stated probabilities. Except near the endpoints, an increase of .05 in the probability of winning increases the value of the prospect by less than 5% of the value of the prize. We next investigate the implications of these psychophysical hypotheses for preferences among risky options.

Figure 1.2. A hypothetical weighting function.

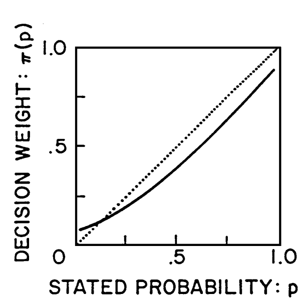

In Figure 1.2, decision weights are lower than the corresponding probabilities over most of the range. Underweighting of moderate and high probabilities relative to sure things contributes to risk aversion in gains by reducing the attractiveness of positive gambles. The same effect also contributes to risk seeking in losses by attenuating the aversiveness of negative gambles. Low probabilities, however, are overweighted, and very low probabilities are either overweighted quite grossly or neglected altogether, making the decision weights highly unstable in that region. The overweighting of low probabilities reverses the pattern described above: It enhances the value of long shots and amplifies the aversiveness of a small chance of a severe loss. Consequently, people are often risk seeking in dealing with improbable gains and risk averse in dealing with unlikely losses. Thus, the characteristics of decision weights contribute to the attractiveness of both lottery tickets and insurance policies.

The nonlinearity of decision weights inevitably leads to violations of invariance, as illustrated in the following pair of problems:

Problem 5 ($N = 85$): Consider the following two-stage game. In the first stage, there is a 75% chance to end the game without winning anything and a 25% chance to move into the second stage. If you reach the second stage you have a choice between:

A. a sure win of $30 (74%)
B. 80% chance to win $45 (26%)

Your choice must be made before the game starts, i.e., before the outcome of the first stage is known. Please indicate the option you prefer.

Problem 6 ($N = 81$): Which of the following options do you prefer?

C. 25% chance to win $30 (42%)
D. 20% chance to win $45 (58%)

Because there is one chance in four to move into the second stage in Problem 5, prospect A offers a .25 probability of winning $30, and prospect B offers $.25 \times .80 = .20$ probability of winning $45. Problems 5 and 6 are therefore identical in terms of probabilities and outcomes. However, the preferences are not the same in the two versions: A clear majority favors the higher chance to win the smaller amount in Problem 5, whereas the majority goes the other way in Problem 6. This violation of invariance has been confirmed with both real and hypothetical monetary payoffs (the present results are with real money), with human lives as outcomes, and with a nonsequential representation of the chance process.

We attribute the failure of invariance to the interaction of two factors: the framing of probabilities and the nonlinearity of decision weights. More specifically, we propose that in Problem 5 people ignore the first phase, which yields the same outcome regardless of the decision that is made, and focus their attention on what happens if they do reach the second stage of the game. In that case, of course, they face a sure gain if they choose option A and an 80% chance of winning if they prefer to gamble. Indeed, people's choices in the sequential version are practically identical to the choices they make between a sure gain of $30 and an 85% chance to win $45. Because a sure thing is overweighted in comparison with events of moderate or high probability (see Figure 1.2) the option that may lead to a gain of $30 is more attractive in the sequential version. We call this phenomenon the *pseudo-certainty* effect because an event that is actually uncertain is weighted as if it were certain.

A closely related phenomenon can be demonstrated at the low end of the probability range. Suppose you are undecided whether or not to purchase earthquake insurance because the premium is quite high. As you hesitate, your friendly insurance agent comes forth with an alternative offer: "For half the regular premium you can be fully covered if the quake occurs on an odd day of the month. This is a good deal because for half the price you are covered for more than half the days." Why do most people find such probabilistic insurance distinctly unattractive? Figure 1.2 suggests an answer. Starting anywhere in the region of low probabilities, the impact on the decision weight of a reduction of probability from p to $p/2$ is considerably smaller than the effect of a reduction from $p/2$ to 0. Reducing the risk by half, then, is not worth half the premium.

The aversion to probabilistic insurance is significant for three reasons. First, it undermines the classical explanation of insurance in terms of a concave utility function. According to expected utility theory, probabilistic insurance should be definitely preferred to normal insurance when the latter is just acceptable (see Kahneman & Tversky, 1979). Second, probabilistic insurance represents many forms of protective action, such as having a medical checkup, buying new tires, or installing a burglar alarm system. Such actions typically reduce the probability of some hazard without eliminating it altogether. Third, the acceptability of insurance can be manipulated by the framing of the contingencies. An insurance policy that covers fire but not flood, for example, could be evaluated either as full protection against a specific risk, (e.g., fire) or as a reduction in the overall probability of property loss. Figure 1.2 suggests that people greatly undervalue a reduction in the probability of a hazard in comparison to the complete elimination of that hazard. Hence, insurance should appear more attractive when it is framed as the elimination of risk than when it is described as a reduction of risk. Indeed, Slovic, Fischhoff, and Lichtenstein (1982) showed that a hypothetical vaccine that reduces the probability of contracting a disease from 20% to 10% is less attractive if it is described as effective in half of the cases than if it is presented as fully effective against one of two exclusive and equally probable virus strains that produce identical symptoms.

Formulation Effects

So far we have discussed framing as a tool to demonstrate failures of invariance. We now turn attention to the processes that control the framing of outcomes and events. The public health problem illustrates a formulation effect in which a change of wording from "lives saved" to "lives lost" induced a marked shift of preference from risk aversion to risk seeking. Evidently, the subjects adopted the descriptions of the outcomes as given in the question and evaluated the outcomes accordingly as gains or losses. Another formulation effect was reported by McNeil, Pauker, Sox, and Tversky (1982). They found that preferences of physicians and patients between hypothetical therapies for lung cancer varied markedly when their probable outcomes were described in terms of mortality or survival. Surgery, unlike radiation therapy, entails a risk of death during treatment. As a consequence, the surgery option was relatively less attractive when the statistics of treatment outcomes were described in terms of mortality rather than in terms of survival.

A physician, and perhaps a presidential advisor as well, could influence the decision made by the patient or by the President, without distorting or suppressing information, merely by the framing of outcomes and contingencies. Formulation effects can occur fortuitously, without anyone being aware of the impact of the frame on the ultimate decision. They can also be exploited deliberately to manipulate the relative attractiveness of options. For example, Thaler (1980) noted that lobbyists for the credit card industry insisted that any price difference between cash and credit purchases be labeled a cash discount rather than a credit card surcharge. The two labels frame the price difference as a gain or as a loss by implicitly designating either the lower or the higher price as normal. Because losses loom larger than gains, consumers are less likely to accept a surcharge than to forego a discount. As is to be expected, attempts to influence framing are common in the marketplace and in the political arena.

The evaluation of outcomes is susceptible to formulation effects because of the nonlinearity of the value function and the tendency of people to evaluate options in relation to the reference point that is suggested or implied by the statement of the problem. It is worthy of note that in other contexts people automatically transform equivalent messages into the same representation. Studies of language comprehension indicate that people quickly recode much of what they hear into an abstract representation that no longer distinguishes whether the idea was expressed in an active or in a passive form and no longer discriminates what was actually said from what was implied, presupposed, or implicated (Clark & Clark, 1977). Unfortunately, the mental machinery that performs these operations silently and effortlessly is not adequate to perform the task of recoding the two versions of the public health problem or the mortality-survival statistics into a common abstract form.

TRANSACTIONS AND TRADES

Our analysis of framing and of value can be extended to choices between multiattribute options, such as the acceptability of a transaction or a trade. We propose that, in order to evaluate a multiattribute option, a person sets up a mental account that specifies the advantages and the disadvantages associated with the option, relative to a multiattribute reference state. The overall value of an option is given by the balance of its advantages and its disadvantages in relation to the reference state. Thus, an option is acceptable if the value of its advantages exceeds the value of its disadvantages. This analysis assumes psychological – but not physical – separability of advantages and disadvantages. The model does not constrain the manner in which separate attributes are combined to form overall measures of advantage and of disadvantage, but it imposes on these measures assumptions of concavity and of loss aversion.

Our analysis of mental accounting owes a large debt to the stimulating work of Richard Thaler (1980, in press), who showed the relevance of this process to consumer behavior. The following problem, based on examples of Savage (1954) and Thaler (1980), introduces some of the rules that govern the construction of mental accounts and illustrates the extension of the concavity of value to the acceptability of transactions.

> **Problem 7:** Imagine that you are about to purchase a jacket for $125 and a calculator for $15. The calculator salesman informs you that the calculator you wish to buy is on sale for $10 at the other branch of the store, located 20 minutes drive away. Would you make a trip to the other store?

This problem is concerned with the acceptability of an option that combines a disadvantage of inconvenience with a financial advantage that can be framed as a *minimal, topical*, or *comprehensive* account. The minimal account includes only the differences between the two options and disregards the features that they share. In the minimal account, the advantage associated with driving to the other store is framed as a gain of $5. A topical account relates the consequences of possible choices to a reference level that is determined by the context within which the decision arises. In the preceding problem, the relevant topic is the purchase of the calculator, and the benefit of the trip is therefore framed as a reduction of the price, from $15 to $10. Because the potential savings is associated only with the calculator, the price of the jacket is not included in the topical account. The price of the jacket, as well as other expenses, could well be included in a more comprehensive account in which the saving would be evaluated in relation to, say, monthly expenses.

The formulation of the preceding problem appears neutral with respect to the adoption of a minimal, topical, or comprehensive account. We suggest, however, that people will spontaneously frame decisions in terms of topical accounts that, in the context of decision making, play a role analogous to that of "good forms" in perception and of basic-level categories in cognition. Topical organization, in conjunction with the concavity of value, entails that the willingness to travel to the other store for a saving of $5 on a calculator should be inversely related to the price of the calculator and should be independent of the price of the jacket. To test this prediction, we constructed another version of the problem in which the prices of the two items were interchanged. The price of the calculator was given as $125 in the first store and $120 in the other branch, and the price of the jacket was set at $15. As predicted, the proportions of respondents who said they would make the trip differed sharply in the two problems. The results showed that 68% of the respondents ($N = 88$) were willing to drive to the other branch to save \$5 on a \$15 calculator, but only 29% of 93 respondents were willing to make the same trip to save \$5 on a \$125 calculator. This finding supports the notion of topical organization of accounts, since the two versions are identical both in terms of a minimal and a comprehensive account.

The significance of topical accounts for consumer behavior is confirmed by the observation that the standard deviation of the prices that different stores in a city quote for the same product is roughly proportional to the average price of that product (Pratt, Wise, & Zeckhauser, 1979). Since the dispersion of prices is surely controlled by shoppers' efforts to find the best buy, these results suggest that consumers hardly exert more effort to save \$15 on a \$150 purchase than to save \$5 on a \$50 purchase.

The topical organization of mental accounts leads people to evaluate gains and losses in relative rather than in absolute terms, resulting in large variations in the rate at which money is exchanged for other things, such as the number of phone calls made to find a good buy or the willingness to drive a long distance to get one. Most consumers will find it easier to buy a car stereo system or a Persian rug, respectively, in the context of buying a car or a house than separately. These observations, of course, run counter to the standard rational theory of consumer behavior, which assumes invariance and does not recognize the effects of mental accounting.

The following problems illustrate another example of mental accounting in which the posting of a cost to an account is controlled by topical organization:

> **Problem 8 ($N = 200$):** Imagine that you have decided to see a play and paid the admission price of \$10 per ticket. As you enter the theater, you discover that you have lost the ticket. The seat was not marked, and the ticket cannot be recovered.
>
> Would you pay \$10 for another ticket?
> Yes (46%)   No (54%)

> **Problem 9 ($N = 183$):** Imagine that you have decided to see a play where admission is \$10 per ticket. As you enter the theater, you discover that you have lost a \$10 bill.
>
> Would you still pay \$10 for a ticket for the play?
> Yes (88%)   No (12%)

The difference between the responses to the two problems is intriguing. Why are so many people unwilling to spend \$10 after having lost a ticket, if they would readily spend that sum after losing an equivalent amount of cash? We attribute the difference to the topical organization of mental accounts. Going to the theater is normally viewed as a transaction in which the cost of the ticket is exchanged for the experience of seeing the play. Buying a second ticket increases the cost of seeing the play to a level that many respondents apparently find unacceptable. In contrast, the loss of the cash is not posted to the account of the play, and it affects the purchase of a ticket only by making the individual feel slightly less affluent.

An interesting effect was observed when the two versions of the problem were presented to the same subjects. The willingness to replace a lost ticket increased significantly when that problem followed the lost-cash version. In contrast, the willingness to buy a ticket after losing cash was not affected by prior presentation of the other problem. The juxtaposition of the two problems apparently enabled the subjects to realize that it makes sense to think of the lost ticket as lost cash, but not vice versa.

The normative status of the effects of mental accounting is questionable. Unlike earlier examples, such as the public health problem, in which the two versions differed only in form, it can be argued that the alternative versions of the calculator and ticket problems differ also in substance. In particular, it may be more pleasurable to save $5 on a $15 purchase than on a larger purchase, and it may be more annoying to pay twice for the same ticket than to lose $10 in cash. Regret, frustration, and self-satisfaction can also be affected by framing (Kahneman & Tversky, 1982). If such secondary consequences are considered legitimate, then the observed preferences do not violate the criterion of invariance and cannot readily be ruled out as inconsistent or erroneous. On the other hand, secondary consequences may change upon reflection. The satisfaction of saving $5 on a $15 item can be marred if the consumer discovers that she would not have exerted the same effort to save $10 on a $200 purchase. We do not wish to recommend that any two decision problems that have the same primary consequences should be resolved in the same way. We propose, however, that systematic examination of alternative framings offers a useful reflective device that can help decision makers assess the values that should be attached to the primary and secondary consequences of their choices.

Losses and Costs

Many decision problems take the form of a choice between retaining the status quo and accepting an alternative to it, which is advantageous in some respects and disadvantageous in others. The analysis of value that was applied earlier to unidimensional risky prospects can be extended to this case by assuming that the status quo defines the reference level for all attributes. The advantages of alternative options will then be evaluated as gains and their disadvantages as losses. Because losses loom larger than gains, the decision maker will be biased in favor of retaining the status quo.

Thaler (1980) coined the term "endowment effect" to describe the reluctance of people to part from assets that belong to their endowment. When it is more painful to give up an asset than it is pleasurable to obtain it, buying prices will be significantly lower than selling prices. That is, the highest price that an individual will pay to acquire an asset will be smaller than the minimal compensation that would induce the same individual to give up that asset, once acquired. Thaler discussed some examples of the endowment effect in the behavior of consumers and entrepreneurs. Several studies have reported substantial discrepancies between buying and selling prices in both hypothetical and real transactions (Gregory, 1983; Hammack & Brown, 1974; Knetsch & Sinden, in press). These results have been presented as challenges to standard economic theory, in which buying and selling prices coincide except for transaction costs and effects of wealth. We also observed reluctance to trade in a study of choices between hypothetical jobs that differed in weekly salary (S) and in the temperature (T) of the workplace. Our respondents were asked to imagine that they held a particular position (S_1, T_1) and were offered the option of moving to a different position (S_2, T_2), which was better in one respect and worse in another. We found that most subjects who were assigned to (S_1, T_1) did not wish to move to (S_2, T_2), and that most subjects who were assigned to the latter position did not wish to move to the former. Evidently, the same difference in pay or in working conditions looms larger as a disadvantage than as an advantage.

In general, loss aversion favors stability over change. Imagine two hedonically identical twins who find two alternative environments equally attractive. Imagine further that by force of circumstance the twins are separated and placed in the two environments. As soon as they adopt their new states as reference points and evaluate the advantages and disadvantages of each other's environments accordingly, the twins will no longer be indifferent between the two states, and both will prefer to stay where they happen to be. Thus, the instability of preferences produces a preference for stability. In addition to favoring stability over change, the combination of adaptation and loss aversion provides limited protection against regret and envy by reducing the attractiveness of foregone alternatives and of others' endowments.

Loss aversion and the consequent endowment effect are unlikely to play a significant role in routine economic exchanges. The owner of a store, for example, does not experience money paid to suppliers as losses and money received from customers as gains. Instead, the merchant adds costs and revenues over some period of time and only evaluates the balance. Matching debits and credits are effectively canceled prior to evaluation. Payments made by consumers are also not evaluated as losses but as alternative purchases. In accord with standard economic analysis, money is naturally viewed as a proxy for the goods and services that it could buy. This mode of evaluation is made explicit when an individual has in mind a particular alternative, such as "I can either buy a new camera or a new tent." In this analysis, a person will buy a camera if its subjective value exceeds the value of retaining the money it would cost.

There are cases in which a disadvantage can be framed either as a cost or as a loss. In particular, the purchase of insurance can also be framed as a choice between a sure loss and the risk of a greater loss. In such cases the cost–loss discrepancy can lead to failures of invariance. Consider, for example, the choice between a sure loss of $50 and a 25% chance to lose $200. Slovic, Fischhoff, and Lichtenstein (1982) reported that 80% of their subjects expressed a risk-seeking preference for the gamble over the sure loss. However, only 35% of subjects refused to pay $50 for insurance against a 25% risk of losing $200. Similar results were also reported by Schoemaker and Kunreuther (1979) and by Hershey and Schoemaker (1980). We suggest that the same amount of money that was framed as an uncompensated loss in the first problem was framed as the cost of protection in the second. The modal preference was reversed in the two problems because losses are more aversive than costs.

We have observed a similar effect in the positive domain, as illustrated by the following pair of problems:

> **Problem 10:** Would you accept a gamble that offers a 10% chance to win $95 and a 90% chance to lose $5?

> **Problem 11:** Would you pay $5 to participate in a lottery that offers a 10% chance to win $100 and a 90% chance to win nothing?

A total of 132 undergraduates answered the two questions, which were separated by a short filler problem. The order of the questions was reversed for half the respondents. Although it is easily confirmed that the two problems offer objectively identical options, 55 of the respondents expressed different preferences in the two versions. Among them, 42 rejected the gamble in Problem 10 but accepted the equivalent lottery in Problem 11. The effectiveness of this seemingly inconsequential manipulation illustrates both the cost-loss discrepancy and the power of framing. Thinking of the $5 as a payment makes the venture more acceptable than thinking of the same amount as a loss.

The preceding analysis implies that an individual's subjective state can be improved by framing negative outcomes as costs rather than as losses. The possibility of such psychological manipulations may explain a paradoxical form of behavior that could be labeled the *dead-loss effect*. Thaler (1980) discussed the example of a man who develops tennis elbow soon after paying the membership fee in a tennis club and continues to play in agony to avoid wasting his investment. Assuming that the individual would not play if he had not paid the membership fee, the question arises: How can playing in agony improve the individual's lot? Playing in pain, we suggest, maintains the evaluation of the membership fee as a cost. If the individual were to stop playing, he would be forced to recognize the fee as a dead loss, which may be more aversive than playing in pain.

CONCLUDING REMARKS

The concepts of utility and value are commonly used in two distinct senses: (a) *experience value*, the degree of pleasure or pain, satisfaction or anguish in the actual experience of an outcome; and (b) *decision value*, the contribution of an anticipated outcome to the overall attractiveness or aversiveness of an option in a choice. The distinction is rarely explicit in decision theory because it is tacitly assumed that decision values and experience values coincide. This assumption is part of the conception of an idealized decision maker who is able to predict future experiences with perfect accuracy and evaluate options accordingly. For ordinary decision makers, however, the correspondence of decision values between experience values is far from perfect (March, 1978). Some factors that affect experience are not easily anticipated, and some factors that affect decisions do not have a comparable impact on the experience of outcomes.

In contrast to the large amount of research on decision making, there has been relatively little systematic exploration of the psychophysics that relate hedonic experience to objective states. The most basic problem of hedonic psychophysics is the determination of the level of adaptation or aspiration that separates positive from negative outcomes. The hedonic reference point is largely determined by the objective status quo, but it is also affected by expectations and social comparisons. An objective improvement can be experienced as a loss, for example, when an employee receives a smaller raise than everyone else in the office. The experience of pleasure or pain associated with a change of state is also critically dependent on the dynamics of hedonic adaptation. Brickman & Campbell's (1971) concept of the hedonic treadmill suggests the radical hypothesis that rapid adaptation will cause the effects of any objective improvement to be short-lived. The complexity and subtlety of hedonic experience make it difficult for the decision maker to anticipate the actual experience that outcomes will produce. Many a person who ordered a meal when ravenously hungry has admitted to a big mistake when the fifth course arrived on the table. The common mismatch of decision values and experience values introduces an additional element of uncertainty in many decision problems.

The prevalence of framing effects and violations of invariance further complicates the relation between decision values and experience values. The framing of outcomes often induces decision values that have no counterpart in actual experience. For example, the framing of outcomes of therapies for lung cancer in terms of mortality or survival is unlikely to affect experience, although it can have a pronounced influence on choice. In other cases, however, the framing of decisions affects not only decision but experience as well. For example, the framing of an expenditure as an uncompensated loss or as the price of insurance can probably influence the experience of that outcome. In such cases, the evaluation of outcomes in the context of decisions not only anticipates experience but also molds it.

2. Prospect Theory

An Analysis of Decision under Risk

Daniel Kahneman and Amos Tversky

ABSTRACT. This paper presents a critique of expected utility theory as a descriptive model of decision making under risk and develops an alternative model, called prospect theory. Choices among risky prospects exhibit several pervasive effects that are inconsistent with the basic tenets of utility theory. In particular, people underweight outcomes that are merely probable in comparison with outcomes that are obtained with certainty. This tendency, called the certainty effect, contributes to risk aversion in choices involving sure gains and to risk seeking in choices involving sure losses. In addition, people generally discard components that are shared by all prospects under consideration. This tendency, called the isolation effect, leads to inconsistent preferences when the same choice is presented in different forms. An alternative theory of choice is developed, in which value is assigned to gains and losses rather than to final assets and in which probabilities are replaced by decision weights. The value function is normally concave for gains, commonly convex for losses, and is generally steeper for losses than for gains. Decision weights are generally lower than the corresponding probabilities, except in the range of low probabilities. Overweighting of low probabilities may contribute to the attractiveness of both insurance and gambling.

1. INTRODUCTION

Expected utility theory has dominated the analysis of decision making under risk. It has been generally accepted as a normative model of rational choice (Keeney & Raiffa, 1976), and widely applied as a descriptive model of economic behavior, e.g. (Friedman & Savage, 1948; Arrow, 1971). Thus, it is assumed that all reasonable people would wish to obey the axioms of the theory (von Neumann & Morgenstern, 1944; Savage, 1954), and that most people actually do, most of the time.

This work was supported in part by grants from the Harry F. Guggenheim Foundation and from the Advanced Research Projects Agency of the Department of Defense and was monitored by Office of Naval Research under Contract N00014-78-C-0100 (ARPA Order No. 3469) under Subcontract 78-072-0722 from Decisions and Designs, Inc. to Perceptronics, Inc. We also thank the Center for Advanced Study in the Behavioral Sciences at Stanford for its support.

Originally published in *Econometrica*, 47:2, 263–91.

The present paper describes several classes of choice problems in which preferences systematically violate the axioms of expected utility theory. In the light of these observations we argue that utility theory, as it is commonly interpreted and applied, is not an adequate descriptive model, and we propose an alternative account of choice under risk.

2. CRITIQUE

Decision making under risk can be viewed as a choice between prospects or gambles. A prospect $(x_1, p_1; \ldots; x_n, p_n)$ is a contract that yields outcome x_i with probability p_i, where $p_1 + p_2 + \cdots + p_n = 1$. To simplify notation, we omit null outcomes and use (x, p) to denote the prospect $(x, p; 0, 1 - p)$ that yields x with probability p and 0 with probability $1 - p$. The (riskless) prospect that yields x with certainty is denoted by (x). The present discussion is restricted to prospects with so-called objective or standard probabilities.

The application of expected utility theory to choices between prospects is based on the following three tenets.

(i) Expectation: $U(x_1, p_1; \ldots; x_n, p_n) = p_1 u(x) + \cdots + p_n u(x_n)$.

That is, the overall utility of a prospect, denoted by U, is the expected utility of its outcomes.

(ii) Asset Integration: $(x_1, p_1; \ldots; x_n, p_n)$ is acceptable at asset position w iff $U(w + x_1, p_1; \ldots; w + x_n, p_n) > u(w)$.

That is, a prospect is acceptable if the utility resulting from integrating the prospect with one's assets exceeds the utility of those assets alone. Thus, the domain of the utility function is final states (which include one's asset position) rather than gains or losses.

Although the domain of the utility function is not limited to any particular class of consequences, most applications of the theory have been concerned with monetary outcomes. Furthermore, most economic applications introduce the following additional assumption.

(iii) Risk Aversion: u is concave $(u'' < 0)$.

A person is risk averse if he prefers the certain prospect (x) to any risky prospect with expected value x. In expected utility theory, risk aversion is equivalent to the concavity of the utility function. The prevalence of risk aversion is perhaps the best known generalization regarding risky choices. It led the early decision theorists of the eighteenth century to propose that utility is a concave function of money, and this idea has been retained in modern treatments (Pratt, 1964; Arrow, 1971).

In the following sections we demonstrate several phenomena which violate these tenets of expected utility theory. The demonstrations are based on the responses of students and university faculty to hypothetical choice problems. The respondents were presented with problems of the type illustrated below.

Which of the following would you prefer?

A: 50% chance to win 1,000, *B: 450 for sure.*
50% chance to win nothing;

The outcomes refer to Israeli currency. To appreciate the significance of the amounts involved, note that the median net monthly income for a family is about 3,000 Israeli pounds. The respondents were asked to imagine that they were actually faced with the choice described in the problem, and to indicate the decision they would have made in such a case. The responses were anonymous, and the instructions specified that there was no 'correct' answer to such problems, and that the aim of the study was to find out how people choose among risky prospects. The problems were presented in questionnaire form, with at most a dozen problems per booklet. Several forms of each questionnaire were constructed so that subjects were exposed to the problems in different orders. In addition, two versions of each problem were used in which the left–right position of the prospects was reversed.

The problems described in this paper are selected illustrations of a series of effects. Every effect has been observed in several problems with different outcomes and probabilities. Some of the problems have also been presented to groups of students and faculty at the University of Stockholm and at the University of Michigan. The pattern of results was essentially identical to the results obtained from Israeli subjects.

The reliance on hypothetical choices raises obvious questions regarding the validity of the method and the generalizability of the results. We are keenly aware of these problems. However, all other methods that have been used to test utility theory also suffer from severe drawbacks. Real choices can be investigated either in the field, by naturalistic or statistical observations of economic behavior, or in the laboratory. Field studies can only provide for rather crude tests of qualitative predictions, because probabilities and utilities cannot be adequately measured in such contexts. Laboratory experiments have been designed to obtain precise measures of utility and probability from actual choices, but these experimental studies typically involve contrived gambles for small stakes, and a large number of repetitions of very similar problems. These features of laboratory gambling complicate the interpretation of the results and restrict their generality.

By default, the method of hypothetical choices emerges as the simplest procedure by which a large number of theoretical questions can be investigated. The use of the method relies on the assumption that people often know how they would behave in actual situations of choice, and on the further assumption that the subjects have no special reason to disguise their true preferences. If people are reasonably accurate in predicting their choices, the presence of common and systematic violations of expected utility theory in hypothetical problems provides presumptive evidence against that theory.

Certainty, Probability, and Possibility

In expected utility theory, the utilities of outcomes are weighted by their probabilities. The present section describes a series of choice problems in which people's preferences systematically violate this principle. We first show that people overweight outcomes that are considered certain, relative to outcomes which are merely probable – a phenomenon which we label the *certainty effect*.

The best known counter-example to expected utility theory which exploits the certainty effect was introduced by the French economist Maurice Allais in 1953. Allais' example has been discussed from both normative and descriptive standpoints by many authors (MacCrimmon and Larsson, forthcoming; Slovic & Tversky, 1974). The following pair of choice problems is a variation of Allais' example, which differs from the original in that it refers to moderate rather than to extremely large gains. The number of respondents who answered each problem is denoted by N, and the percentage who choose each option is given in brackets.

Problem 1: Choose between

A: 2,500 with probability	.33,	B: 2,400 with certainty.		
2,400 with probability	.66,			
0 with probability	.01;			
$N = 72$ [18]		[82]*		

Problem 2: Choose between

C: 2,500 with probability	.33,	D: 2,400 with probability	.34,
0 with probability	.67;	0 with probability	.66.
$N = 72$ [83]*		[17]	

The data show that 82 percent of the subjects chose B in Problem 1, and 83 percent of the subjects chose C in Problem 2. Each of these preferences is significant at the .01 level, as denoted by the asterisk. Moreover, the analysis of individual patterns of choice indicates that a majority of respondents (61 percent) made the modal choice in both problems. This pattern of preferences violates expected utility theory in the manner originally described by Allais. According to that theory, with $u(0) = 0$, the first preference implies

$$u(2{,}400) > .33u(2{,}500) + .66u(2{,}400) \quad \text{or} \quad .34u(2{,}400) > .33u(2{,}500)$$

while the second preference implies the reverse inequality. Note that Problem 2 is obtained from Problem 1 by eliminating a .66 chance of winning 2,400 from both prospects under consideration. Evidently, this change produces a greater reduction in desirability when it alters the character of the prospect from a sure gain to a probable one, than when both the original and the reduced prospects are uncertain.

A simpler demonstration of the same phenomenon, involving only two-outcome gambles is given below. This example is also based on Allais (1953).

Problem 3:

A: (4,000, .80), or B: (3,000).
$N = 95$ [20] [80]*

Problem 4:

C: (4,000, .20), or D: (3,000, .25).
$N = 95$ [65]* [35]

In this pair of problems as well as in all other problem-pairs in this section, over half the respondents violated expected utility theory. To show that the modal pattern of preferences in Problems 3 and 4 is not compatible with the theory, set $u(0) = 0$, and recall that the choice of B implies $u(3{,}000)/u(4{,}000) > 4/5$, whereas the choice of C implies the reverse inequality. Note that the prospect $C = (4{,}000, .20)$ can be expressed as $(A, .25)$, while the prospect $D = (3{,}000, .25)$ can be rewritten as $(B, .25)$. The substitution axiom of utility theory asserts that if B is preferred to A, then any (probability) mixture (B, p) must be preferred to the mixture (A, p). Our subjects did not obey this axiom. Apparently, reducing the probability of winning from 1.0 to .25 has a greater effect than the reduction from .8 to .2. The following pair of choice problems illustrates the certainty effect with non-monetary outcomes.

Problem 5:

A: 50% chance to win a three-week tour of England, France, and Italy;
B: A one-week tour of England, with certainty.
$N = 72$ [22] [78]*

Problem 6:

C: 5% chance to win a three-week tour of England, France, and Italy;
D: 10% chance to win a one-week tour of England.
$N = 72$ [67]* [33]

The certainty effect is not the only type of violation of the substitution axiom. Another situation in which this axiom fails is illustrated by the following problems.

Problem 7:

A: (6,000, .45), B: (3,000, .90).
$N = 66$ [14] [86]*

Problem 8:

C: (6,000, .001),	D: (3,000, .002).
$N = 66$ [73]*	[27]

Note that in Problem 7 the probabilities of winning are substantial (.90 and .45), and most people choose the prospect where winning is more probable. In Problem 8, there is a *possibility* of winning, although the probabilities of winning are minuscule (.002 and .001) in both prospects. In this situation where winning is possible but not probable, most people choose the prospect that offers the larger gain. Similar results have been reported by MacCrimmon and Larsson (forthcoming).

The above problems illustrate common attitudes toward risk or chance that cannot be captured by the expected utility model. The results suggest the following empirical generalization concerning the manner in which the substitution axiom is violated. If (y, pq) is equivalent to (x, p), then (y, pqr) is preferred to (x, pr), $0 < p, q, r < 1$. This property is incorporated into an alternative theory, developed in the second part of the paper.

The Reflection Effect

The previous section discussed preferences between positive prospects, i.e., prospects that involve no losses. What happens when the signs of the outcomes are reversed so that gains are replaced by losses? The left-hand column of Table 2.1 displays four of the choice problems that were discussed in the previous section, and the right-hand column displays choice problems in which the signs of the outcomes are reversed. We use $-x$ to denote the loss of x, and $>$ to denote the prevalent preference, i.e., the choice made by the majority of subjects.

In each of the four problems in Table 2.1 the preference between negative prospects is the mirror image of the preference between positive prospects. Thus, the reflection of prospects around 0 reverses the preference order. We label this pattern the *reflection effect*.

Let us turn now to the implications of these data. First, note that the reflection effect implies that risk aversion in the positive domain is accompanied by

Table 2.1. Preferences between Positive and Negative Prospects

Positive Prospects				Negative Prospects			
Problem 3:	(4,000, .80)	<	(3,000).	Problem 3′:	(−4,000, .80)	>	(−3,000).
$N = 95$	[20]		[80]*	$N = 95$	[92]*		[8]
Problem 4:	(4,000, .20)	>	(3,000, .25).	Problem 4′:	(−4,000, .20)	<	(−3,000, .25).
$N = 95$	[65]*		[35]	$N = 95$	[42]		[58]
Problem 7:	(3,000, .90)	>	(6,000, .45).	Problem 7′:	(−3,000, .90)	<	(−6,000, .45).
$N = 66$	[86]*		[14]	$N = 66$	[8]		[92]*
Problem 8:	(3,000, .002)	<	(6,000, .001).	Problem 8′:	(−3,000, .002)	>	(−6,000, .001).
$N = 66$	[27]		[73]*	$N = 66$	[70]*		[30]

risk seeking in the negative domain. In Problem 3′, for example, the majority of subjects were willing to accept a risk of .80 to lose 4,000, in preference to a sure loss of 3,000, although the gamble has a lower expected value. The occurrence of risk seeking in choices between negative prospects was noted early by Markowitz (1952). Williams (1966) reported data where a translation of outcomes produces a dramatic shift from risk aversion to risk seeking. For example, his subjects were indifferent between (100, .65; −100, .35) and (0), indicating risk aversion. They were also indifferent between (−200, .80) and (−100), indicating risk seeking. A recent review by Fishburn and Kochenberger (1979) documents the prevalence of risk seeking in choices between negative prospects.

Second, recall that the preferences between the positive prospects in Table 2.1 are inconsistent with expected utility theory. The preferences between the corresponding negative prospects also violate the expectation principle in the same manner. For example, Problems 3′ and 4′, like Problems 3 and 4, demonstrate that outcomes which are obtained with certainty are overweighted relative to uncertain outcomes. In the positive domain, the certainty effect contributes to a risk averse preference for a sure gain over a larger gain that is merely probable. In the negative domain, the same effect leads to a risk-seeking preference for a loss that is merely probable over a smaller loss that is certain. The same psychological principle – the overweighting of certainty – favors risk aversion in the domain of gains and risk seeking in the domain of losses.

Third, the reflection effect eliminates aversion for uncertainty or variability as an explanation of the certainty effect. Consider, for example, the prevalent preferences for (3,000) over (4,000, .80) and for (4,000, .20) over (3,000, .25). To resolve this apparent inconsistency one could invoke the assumption that people prefer prospects that have high expected value and small variance (see, e.g., Allais (1953); Markowitz (1959); Tobin (1958)). Since (3,000) has no variance while (4,000, .80) has large variance, the former prospect could be chosen despite its lower expected value. When the prospects are reduced, however, the difference in variance between (3,000, .25) and (4,000, .20) may be insufficient to overcome the difference in expected value. Because (−3,000) has both higher expected value and lower variance than (−4,000, .80), this account entails that the sure loss should be preferred, contrary to the data. Thus, our data are incompatible with the notion that certainty is generally desirable. Rather, it appears that certainty increases the aversiveness of losses as well as the desirability of gains.

Probabilistic Insurance

The prevalence of the purchase of insurance against both large and small losses has been regarded by many as strong evidence for the concavity of the utility function for money. Why otherwise would people spend so much money to purchase insurance policies at a price that exceeds the expected actuarial cost? However, an examination of the relative attractiveness of various forms of insurance does not support the notion that the utility function for money is concave everywhere. For example, people often prefer insurance programs that offer limited coverage with low or zero deductible over comparable policies that offer higher maximal coverage with higher deductibles – contrary to risk aversion (see, e.g., Fuchs (1976)). Another type of insurance problem in which people's responses are inconsistent with the concavity hypothesis may be called probabilistic insurance. To illustrate this concept, consider the following problem, which was presented to 95 Stanford University students.

Problem 9: Suppose you consider the possibility of insuring some property against damage, e.g., fire or theft. After examining the risks and the premium you find that you have no clear preference between the options of purchasing insurance or leaving the property uninsured.

It is then called to your attention that the insurance company offers a new program called *probabilistic insurance*. In this program you pay half of the regular premium. In case of damage, there is a 50 percent chance that you pay the other half of the premium and the insurance company covers all the losses, and there is a 50 percent chance that you get back your insurance payment and suffer all the losses. For example, if an accident occurs on an odd day of the month, you pay the other half of the regular premium and your losses are covered, but if the accident occurs on an even day of the month, your insurance payment is refunded and your losses are not covered.

Recall that the premium for full coverage is such that you find this insurance barely worth its cost.

Under these circumstances, would you purchase probabilistic insurance:

	Yes,	No.
$N = 95$	[20]	[80]*

Although Problem 9 may appear contrived, it is worth noting that probabilistic insurance represents many forms of protective action where one pays a certain cost to reduce the probability of an undesirable event – without eliminating it altogether. The installation of a burglar alarm, the replacement of old tires, and the decision to stop smoking can all be viewed as probabilistic insurance.

The responses to Problem 9 and to several other variants of the same question indicate that probabilistic insurance is generally unattractive. Apparently, reducing the probability of a loss from p to $p/2$ is less valuable than reducing the probability of that loss from $p/2$ to 0.

In contrast to these data, expected utility theory (with a concave u) implies that probabilistic insurance is superior to regular insurance. That is, if at asset position w one is just willing to pay a premium y to insure against a probability p of losing x, then one should definitely be willing to pay a smaller premium ry to reduce the probability of losing x from p to $(1-r)p$, $0 < r < 1$. Formally, if one is indifferent between $(w-x,\ p; w,\ 1-p)$ and $(w-y)$, then one should prefer probabilistic insurance $(w - x, (1 - r)p; w - y, rp; w - ry, 1 - p)$ over regular insurance $(w - y)$.

To prove this proposition, we show that

$$pu(w - x) + (1 - p)u(w) = u(w - y)$$

implies

$$(1 - r)pu(w - x) + rpu(w - y) + (1 - p)u(w - ry) > u(w - y).$$

Without loss of generality, we can set $u(w - x) = 0$ and $u(w) = 1$. Hence, $u(w - y) = 1 - p$, and we wish to show that

$$rp(1 - p) + (1 - p)u(w - ry) > 1 - p \quad \text{or} \quad u(w - ry) > 1 - rp$$

which holds if and only if u is concave.

This is a rather puzzling consequence of the risk aversion hypothesis of utility theory, because probabilistic insurance appears intuitively riskier than regular insurance, which entirely eliminates the element of risk. Evidently, the intuitive notion of risk is not adequately captured by the assumed concavity of the utility function for wealth.

The aversion for probabilistic insurance is particularly intriguing because all insurance is, in a sense, probabilistic. The most avid buyer of insurance remains vulnerable to many financial and other risks which his policies do not cover. There appears to be a significant difference between probabilistic insurance and what may be called contingent insurance, which provides the certainty of coverage for a specified type of risk. Compare, for example, probabilistic insurance against all forms of loss or damage to the contents of your home and contingent insurance that eliminates all risk of loss from theft, say, but does not cover other risks, e.g., fire. We conjecture that contingent insurance will be generally more attractive than probabilistic insurance when the probabilities of unprotected loss are equated. Thus, two prospects that are equivalent in probabilities and outcomes could have different values depending on their formulation. Several demonstrations of this general phenomenon are described in the next section.

The Isolation Effect

In order to simplify the choice between alternatives, people often disregard components that the alternatives share and focus on the components that distinguish them (Tversky (1972)). This approach to choice problems may produce inconsistent preferences, because a pair of prospects can be decomposed into common and distinctive components in more than one way, and different decompositions sometimes lead to different preferences. We refer to this phenomenon as the *isolation effect*.

Problem 10: Consider the following two-stage game. In the first stage, there is a probability of .75 to end the game without winning anything, and a probability

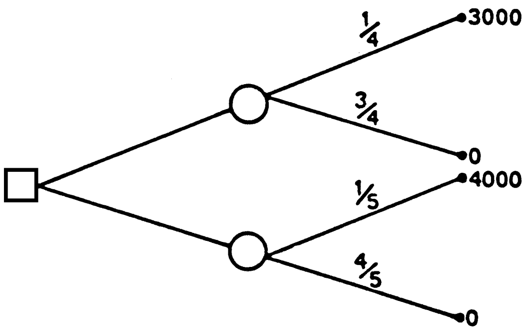

Figure 2.1. The representation of Problem 4 as a decision tree (standard formulation).

of .25 to move into the second stage. If you reach the second stage you have a choice between

(4,000, .80) and (3,000).

Your choice must be made before the game starts, i.e., before the outcome of the first stage is known.

Note that in this game, one has a choice between $.25 \times .80 = .20$ chance to win 4,000, and a $.2 \times 1.0 = .25$ chance to win 3,000. Thus, in terms of final outcomes and probabilities one faces a choice between (4,000, .20) and (3,000, .25), as in Problem 4 above. However, the dominant preferences are different in the two problems. Of 141 subjects who answered Problem 10, 78 percent chose the latter prospect, contrary to the modal preference in Problem 4. Evidently, people ignored the first stage of the game, whose outcomes are shared by both prospects, and considered Problem 10 as a choice between (3,000) and (4,000, .80), as in Problem 3 above.

The standard and the sequential formulations of Problem 4 are represented as decision trees in Figures 2.1 and 2.2, respectively. Following the usual convention, squares denote decision nodes and circles denote chance nodes. The essential difference between the two representations is in the location of the

Figure 2.2. The representation of Problem 10 as a decision tree (sequential formulation).

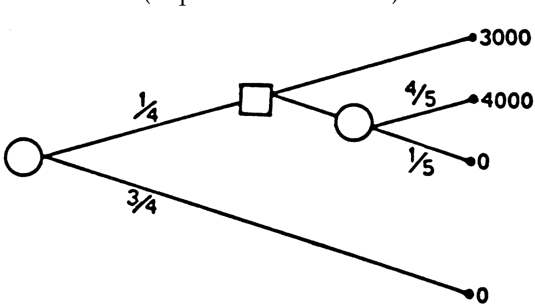

decision node. In the standard form (Figure 2.1), the decision maker faces a choice between two risky prospects, whereas in the sequential form (Figure 2.2) he faces a choice between a risky and a riskless prospect. This is accomplished by introducing a dependency between the prospects without changing either probabilities or outcomes. Specifically, the event 'not winning 3,000' is included in the event 'not winning 4,000' in the sequential formulation, while the two events are independent in the standard formulation. Thus, the outcome of winning 3,000 has a certainty advantage in the sequential formulation, which it does not have in the standard formulation.

The reversal of preferences due to the dependency among events is particularly significant because it violates the basic supposition of a decision-theoretical analysis, that choices between prospects are determined solely by the probabilities of final states.

It is easy to think of decision problems that are most naturally represented in one of the forms above rather than in the other. For example, the choice between two different risky ventures is likely to be viewed in the standard form. On the other hand, the following problem is most likely to be represented in the sequential form. One may invest money in a venture with some probability of losing one's capital if the venture fails, and with a choice between a fixed agreed return and a percentage of earnings if it succeeds. The isolation effect implies that the contingent certainty of the fixed return enhances the attractiveness of this option, relative to a risky venture with the same probabilities and outcomes.

The preceding problem illustrated how preferences may be altered by different representations of probabilities. We now show how choices may be altered by varying the representation of outcomes.

Consider the following problems, which were presented to two different groups of subjects.

Problem 11: In addition to whatever you own, you have been given 1,000. You are now asked to choose between

A: (1,000, .50), and B: (500).
$N = 70$ [16] [84]*

Problem 12: In addition to whatever you own, you have been given 2,000. You are now asked to choose between

C: (−1,000, .50), and D: (−500).
$N = 68$ [69*] [31]

The majority of subjects chose *B* in the first problem and *C* in the second. These preferences conform to the reflection effect observed in Table 2.1, which exhibits risk aversion for positive prospects and risk seeking for negative ones. Note, however, that when viewed in terms of final states, the two choice problems are identical. Specifically,

A = (2,000, .50; 1,000, .50) = C, and B = (1,500) = D.

In fact, Problem 12 is obtained from Problem 11 by adding 1,000 to the initial bonus, and subtracting 1,000 from all outcomes. Evidently, the subjects did not integrate the bonus with the prospects. The bonus did not enter into the comparison of prospects because it was common to both options in each problem.

The pattern of results observed in Problems 11 and 12 is clearly inconsistent with utility theory. In that theory, for example, the same utility is assigned to a wealth of $100,000, regardless of whether it was reached from a prior wealth of $95,000 or $105,000. Consequently, the choice between a total wealth of $100,000 and even chances to own $95,000 or $105,000 should be independent of whether one currently owns the smaller or the larger of these two amounts. With the added assumption of risk aversion, the theory entails that the certainty of owning $100,000 should always be preferred to the gamble. However, the responses to Problem 12 and to several of the previous questions suggest that this pattern will be obtained if the individual owns the smaller amount, but not if he owns the larger amount.

The apparent neglect of a bonus that was common to both options in Problems 11 and 12 implies that the carriers of value or utility are changes of wealth, rather than final asset positions that include current wealth. This conclusion is the cornerstone of an alternative theory of risky choice, which is described in the following sections.

3. THEORY

The preceding discussion reviewed several empirical effects which appear to invalidate expected utility theory as a descriptive model. The remainder of the paper presents an alternative account of individual decision making under risk, called prospect theory. The theory is developed for simple prospects with monetary outcomes and stated probabilities, but it can be extended to more involved choices. Prospect theory distinguishes two phases in the choice process: an early phase of editing and a subsequent phase of evaluation. The editing phase consists of a preliminary analysis of the offered prospects, which often yields a simpler representation of these prospects. In the second phase, the edited prospects are evaluated and the prospect of highest value is chosen. We next outline the editing phase and develop a formal model of the evaluation phase.

The function of the editing phase is to organize and reformulate the options so as to simplify subsequent evaluation and choice. Editing consists of the application of several operations that transform the outcomes and probabilities associated with the offered prospects. The major operations of the editing phase are described below.

Coding: The evidence discussed in the previous section shows that people normally perceive outcomes as gains and losses, rather than as final states of wealth or welfare. Gains and losses, of course, are defined relative to some neutral reference point. The reference point usually corresponds to the current asset position, in which case gains and losses coincide with the actual amounts that are received or paid. However, the location of the reference point, and the consequent coding of outcomes as gains or losses, can be affected by the formulation of the offered prospects and by the expectations of the decision maker.

Combination: Prospects can sometimes be simplified by combining the probabilities associated with identical outcomes. For example, the prospect (200, .25; 200, .25) will be reduced to (200, .50), and evaluated in this form.

Segregation: Some prospects contain a riskless component that is segregated from the risky component in the editing phase. For example, the prospect (300, .80; 200, .20) is naturally decomposed into a sure gain of 200 and the risky prospect (100, .80). Similarly, the prospect (−400, .40; −100, .60) is readily seen to consist of a sure loss of 100 and of the prospect (−300, .40).

The preceding operations are applied to each prospect separately. The following operation is applied to a set of two or more prospects.

Cancellation: The essence of the isolation effects described earlier is the discarding of components that are shared by the offered prospects. Thus, our respondents apparently ignored the first stage of the sequential game presented in Problem 10, because this stage was common to both options, and they evaluated the prospects with respect to the results of the second stage (see Figure 2.2). Similarly, they neglected the common bonus that was added to the prospects in Problems 11 and 12. Another type of cancellation involves the discarding of common constituents, i.e., outcome-probability pairs. For example, the choice between (200, .20; 100, .50; −50, .30) and (200, .20; 150, .50; −100, .30) can be reduced by cancellation to a choice between (100, .50; −50, .30) and (150, .50; −100, .30).

Two additional operations that should be mentioned are simplification and the detection of dominance. The first refers to the simplification of prospects by rounding probabilities or outcomes. For example, the prospect (101, .49) is likely to be recoded as an even chance to win 100. A particularly important form of simplification involves the discarding of extremely unlikely outcomes. The second operation involves the scanning of offered prospects to detect dominated alternatives, which are rejected without further evaluation.

Because the editing operations facilitate the task of decision, it is assumed that they are performed whenever possible. However, some editing operations either permit or prevent the application of others. For example, (500, .20; 101, .49) will appear to dominate (500, .15; 99, .51) if the second constituents of both prospects are simplified to (100, .50). The final edited prospects could, therefore, depend on the sequence of editing operations, which is likely to vary with the structure of the offered set and with the format of the display. A detailed study of this problem is beyond the scope of the present treatment. In this paper we discuss choice problems where it is reasonable to assume either that the original formulation of the prospects leaves no room for further editing, or that the edited prospects can be specified without ambiguity.

Many anomalies of preference result from the editing of prospects. For example, the inconsistencies associated with the isolation effect result from the cancellation of common components. Some intransitivities of choice are explained by a simplification that eliminates small differences between prospects (see Tversky (1969)). More generally, the preference order between prospects need not be invariant across contexts, because the same offered prospect could be edited in different ways depending on the context in which it appears.

Following the editing phase, the decision maker is assumed to evaluate each of the edited prospects, and to choose the prospect of highest value. The overall value of an edited prospect, denoted V, is expressed in terms of two scales, π and v.

The first scale, π, associates with each probability p a decision weight $\pi(p)$, which reflects the impact of p on the overall value of the prospect. However, π is not a probability measure, and it will be shown later that $\pi(p) + \pi(1 - p)$ is typically less than unity. The second scale, v, assigns to each outcome x a number $v(x)$, which reflects the subjective value of that outcome. Recall that outcomes are defined relative to a reference point, which serves as the zero point of the value scale. Hence, v measures the value of deviations from that reference point, i.e., gains and losses.

The present formulation is concerned with simple prospects of the form $(x, p; y, q)$, which have at most two non-zero outcomes. In such a prospect, one receives x with probability p, y with probability q, and nothing with probability $1 - p - q$, where $p + q \leq 1$. An offered prospect is strictly positive if its outcomes are all positive, i.e., if $x, y > 0$ and $p + q = 1$; it is strictly negative if its outcomes are all negative. A prospect is regular if it is neither strictly positive nor strictly negative.

The basic equation of the theory describes the manner in which π and v are combined to determine the overall value of regular prospects.

If $(x, p; y, q)$ is a regular prospect (i.e., either $p + q < 1$, or $x \geq 0 \geq y$, or $x \leq 0 \leq y$), then

$$V(x, p; y, q) = \pi(p)v(x) + \pi(q)v(y), \tag{1}$$

where $v(0) = 0$, $\pi(0) = 0$, and $\pi(1) = 1$. As in utility theory, V is defined on prospects, while v is defined on outcomes. The two scales coincide for sure prospects, where $V(x, 1.0) = V(x) = v(x)$.

Equation (1) generalizes expected utility theory by relaxing the expectation principle. An axiomatic analysis of this representation is sketched in the Appendix (not printed here), which describes conditions that ensure the existence of a unique π and a ratio-scale v satisfying equation (1).

The evaluation of strictly positive and strictly negative prospects follows a different rule. In the editing phase such prospects are segregated into two components: (i) the riskless component, i.e., the minimum gain or loss which is certain to be obtained or paid; (ii) the risky component, i.e., the additional gain or loss which is actually at stake. The evaluation of such prospects is described in the next equation.

If $p + q = 1$ and either $x > y > 0$ or $x < y < 0$, then

$$V(x, p; y, q) = v(y) + \pi(p)[v(x) - v(y)]. \tag{2}$$

That is, the value of a strictly positive or strictly negative prospect equals the value of the riskless component plus the value-difference between the outcomes, multiplied by the weight associated with the more extreme outcome. For example, $V(400, .25; 100, .75) = v(100) + \pi(.25)[v(400) - v(100)]$. The essential feature of equation (2) is that a decision weight is applied to the value-difference $v(x) - v(y)$, which represents the risky component of the prospect, but not to $v(y)$, which represents the riskless component. Note that the right-hand side of equation (2) equals $\pi(p)v(x) + [1 - \pi(p)]v(y)$. Hence, equation (2) reduces to equation (1) if $\pi(p) + \pi(1 - p) = 1$. As will be shown later, this condition is not generally satisfied.

Many elements of the evaluation model have appeared in previous attempts to modify expected utility theory. Markowitz (1952) was the first to propose that utility be defined on gains and losses rather than on final asset positions, an assumption which has been implicitly accepted in most experimental measurements of utility (see, e.g., (Davidson, Suppes & Siegel, 1957; Mosteller & Nogee, 1951)). Markowitz also noted the presence of risk seeking in preferences among positive as well as among negative prospects, and he proposed a utility function which has convex and concave regions in both the positive and the negative domains. His treatment, however, retains the expectation principle; hence it cannot account for the many violations of this principle; see, e.g., Table 2.1.

The replacement of probabilities by more general weights was proposed by Edwards (1962), and this model was investigated in several empirical studies (e.g., (Anderson and Shanteau, 1970; Tversky, 1967)). Similar models were developed by Fellner (1965), who introduced the concept of decision weight to explain aversion for ambiguity, and by van Dam (1975) who attempted to scale decision weights. For other critical analyses of expected utility theory and alternative choice models, see Allais (1953), Coombs (1975), Fishburn (1977), and Hansson (1975).

The equations of prospect theory retain the general bilinear form that underlies expected utility theory. However, in order to accommodate the effects described in the first part of the paper, we are compelled to assume that values are attached to changes rather than to final states, and that decision weights do not coincide with stated probabilities. These departures from expected utility theory must lead to normatively unacceptable consequences, such as inconsistencies, intransitivities, and violations of dominance. Such anomalies of preference are normally corrected by the decision maker when he realizes that his preferences are inconsistent, intransitive, or inadmissible. In many situations, however, the decision maker does not have the opportunity to discover that his preferences could violate decision rules that he wishes to obey. In these circumstances the anomalies implied by prospect theory are expected to occur.

The Value Function

An essential feature of the present theory is that the carriers of value are changes in wealth or welfare, rather than final states. This assumption is compatible with basic principles of perception and judgment. Our perceptual apparatus is attuned to the evaluation of changes or differences rather than to the evaluation of absolute magnitudes. When we respond to attributes such as brightness, loudness, or temperature, the past and present context of experience defines an adaptation level, or reference point, and stimuli are perceived in relation to this reference point (Helson, 1964). Thus, an object at a given temperature may be experienced as hot or cold to the touch depending on the temperature to which one has adapted. The same principle applies to non-sensory attributes such as health, prestige, and wealth. The same level of wealth, for example, may imply abject poverty for one person and great riches for another – depending on their current assets.

The emphasis on changes as the carriers of value should not be taken to imply that the value of a particular change is independent of initial position. Strictly speaking, value should be treated as a function in two arguments: the asset position that serves as reference point, and the magnitude of the change (positive or negative) from that reference point. An individual's attitude to money, say, could be described by a book, where each page presents the value function for changes at a particular asset position. Clearly, the value functions described on different pages are not identical: they are likely to become more linear with increases in assets. However, the preference order of prospects is not greatly altered by small or even moderate variations in asset position. The certainty equivalent of the prospect (1,000, .50), for example, lies between 300 and 400 for most people, in a wide range of asset positions. Consequently, the representation of value as a function in one argument generally provides a satisfactory approximation.

Many sensory and perceptual dimensions share the property that the psychological response is a concave function of the magnitude of physical change. For example, it is easier to discriminate between a change of 3° and a change of 6° in room temperature, than it is to discriminate between a change of 13° and a change of 16°. We propose that this principle applies in particular to the evaluation of monetary changes. Thus, the difference in value between a gain of 100 and a gain of 200 appears to be greater than the difference between a gain of 1,100 and a gain of 1,200. Similarly, the difference between a loss of 100 and a loss of 200 appears greater than the difference between a loss of 1,100 and a loss of 1,200, unless the larger loss is intolerable. Thus, we hypothesize that the value function for changes of wealth is normally concave above the reference point ($v''(x) < 0$, for $x > 0$) and often convex below it ($v''(x) > 0$, for $x < 0$). That is, the marginal value of both gains and losses generally decreases with their magnitude. Some support for this hypothesis has been reported by Galanter and Pliner (1974), who scaled the perceived magnitude of monetary and non-monetary gains and losses.

The above hypothesis regarding the shape of the value function was based on responses to gains and losses in a riskless context. We propose that the value function which is derived from risky choices shares the same characteristics, as illustrated in the following problems.

Problem 13:

	(6,000, .25),	or	(4,000, .25; 2,000, .25).
$N = 68$	[18]		[82]*

Problem 13′:

	(−6,000, .25),	or	(−4,000, .25; −2,000, .25).
$N = 64$	[70]*		[30]

Applying equation 1 to the modal preference in these problems yields

$$\pi(.25)v(6{,}000) < \pi(.25)[v(4{,}000) + v(2{,}000)] \quad \text{and}$$
$$\pi(.25)v(-6{,}000) > \pi(.25)[v(-4{,}000) + v(-2{,}000)].$$

Hence, $v(6{,}000) < v(4{,}000) + v(2{,}000)$ and $v(-6{,}000) > v(-4{,}000) + v(-2{,}000)$. These preferences are in accord with the hypothesis that the value function is concave for gains and convex for losses.

Any discussion of the utility function for money must leave room for the effect of special circumstances on preferences. For example, the utility function of an individual who needs $60,000 to purchase a house may reveal an exceptionally steep rise near the critical value. Similarly, an individual's aversion to losses may increase sharply near the loss that would compel him to sell his house and move to a less desirable neighborhood. Hence, the derived value (utility) function of an individual does not always reflect "pure" attitudes to money, since it could be affected by additional consequences associated with specific amounts. Such perturbations can readily produce convex regions in the value function for gains and concave regions in the value function for losses. The latter case may be more common since large losses often necessitate changes in life style.

A salient characteristic of attitudes to changes in welfare is that losses loom larger than gains. The aggravation that one experiences in losing a sum of money appears to be greater than the pleasure associated with gaining the same amount (Galanter & Pliner, 1974). Indeed, most people find symmetric bets of the form $(x, .50; -x, .50)$ distinctly unattractive. Moreover, the aversiveness of symmetric fair bets generally increases with the size of the stake. That is, if $x > y \geq 0$, then $(y, .50; -y, .50)$ is preferred to $(x, .50; -x, .50)$. According to equation (1), therefore,

$$v(y) + v(-y) > v(x) + v(-x) \quad \text{and} \quad v(-y) - v(-x) > v(x) - v(y).$$

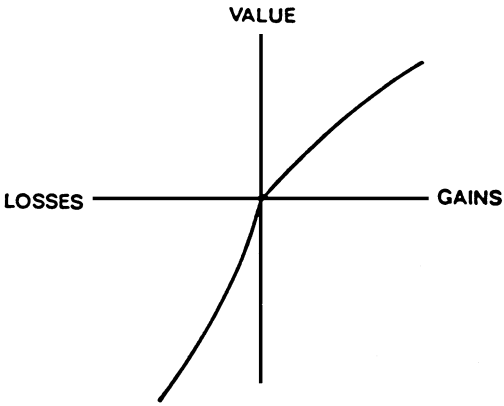

Figure 2.3. A hypothetical value function.

Setting $y=0$ yields $v(x) < -v(-x)$, and letting y approach x yields $v'(x) < v'(-x)$, provided v', the derivative of v, exists. Thus, the value function for losses is steeper than the value function for gains.

In summary, we have proposed that the value function is (i) defined on deviations from the reference point; (ii) generally concave for gains and commonly convex for losses; (iii) steeper for losses than for gains. A value function which satisfies these properties is displayed in Figure 2.3. Note that the proposed S-shaped value function is steepest at the reference point, in marked contrast to the utility function postulated by Markowitz (1952) which is relatively shallow in that region.

Although the present theory can be applied to derive the value function from preferences between prospects, the actual scaling is considerably more complicated than in utility theory because of the introduction of decision weights. For example, decision weights could produce risk aversion and risk seeking even with a linear value function. Nevertheless, it is of interest that the main properties ascribed to the value function have been observed in a detailed analysis of von Neumann–Morgenstern utility functions for changes of wealth (Fishburn and Kochenberger (1979)). The functions had been obtained from thirty decision makers in various fields of business, in five independent studies (Barnes & Reinmuth, 1976; Grayson, 1960; Green, 1963; Halter & Dean, 1971; Swalm, 1966). Most utility functions for gains were concave, most functions for losses were convex, and only three individuals exhibited risk aversion for both gains and losses. With a single exception, utility functions were considerably steeper for losses than for gains.

The Weighting Function

In prospect theory, the value of each outcome is multiplied by a decision weight. Decision weights are inferred from choices between prospects much as subjective probabilities are inferred from preferences in the Ramsey–Savage approach. However, decision weights are not probabilities: they do not obey the probability axioms and they should not be interpreted as measures of degree or belief.

Consider a gamble in which one can win 1,000 or nothing, depending on the toss of a fair coin. For any reasonable person, the probability of winning is .50 in this situation. This can be verified in a variety of ways, e.g., by showing that the subject is indifferent between betting on heads or tails, or by his verbal report that he considers the two events equiprobable. As will be shown below, however, the decision weight $\pi(.50)$ which is derived from choices is likely to be smaller than .50. Decision weights measure the impact of events on the desirability of prospects, and not merely the perceived likelihood of these events. The two scales coincide (i.e., $\pi(p) = p$) if the expectation principle holds, but not otherwise.

The choice problems discussed in the present paper were formulated in terms of explicit numerical probabilities, and our analysis assumes that the respondents adopted the stated values of p. Furthermore, since the events were identified only by their stated probabilities, it is possible in this context to express decision weights as a function of stated probability. In general, however, the decision weight attached to an event could be influenced by other factors, e.g., ambiguity (Ellsberg, 1961; Fellner, 1961).

We turn now to discuss the salient properties of the weighting function π, which relates decision weights to stated probabilities. Naturally, π is an increasing function of p, with $\pi(0) = 0$ and $\pi(1) = 1$. That is, outcomes contingent on an impossible event are ignored, and the scale is normalized so that $\pi(p)$ is the ratio of the weight associated with the probability p to the weight associated with the certain event.

We first discuss some properties of the weighting function for small probabilities. The preferences in Problems 8 and 8′ suggest that for small values of p, π is a subadditive function of p, i.e., $(rp) > r\pi(p)$ for $0 < r < 1$. Recall that in Problem 8, (6,000, .001) is preferred to (3,000, .002). Hence

$$\frac{\pi(.001)}{\pi(.002)} > \frac{v(3{,}000)}{v(6{,}000)} > \frac{1}{2} \qquad \text{by the concavity of } v.$$

The reflected preferences in Problem 8′ yield the same conclusion. The pattern of preferences in Problems 7 and 7′, however, suggests that subadditivity need not hold for large values of p.

Furthermore, we propose that very low probabilities are generally overweighted, that is, $\pi(p) > p$ for small p. Consider the following choice problems.

Problem 14:

	(5,000, .001),	or	(5).
$N = 72$	[72]*		[28]

Problem 14′:

	(−5,000, .001),	or	(−5).
$N = 72$	[17]		[83]*

Note that in Problem 14, people prefer what is in effect a lottery ticket over the expected value of that ticket. In Problem 14′, on the other hand, they prefer a small loss, which can be viewed as the payment of an insurance premium, over a small probability of a large loss. Similar observations have been reported by Markowitz (1952). In the present theory, the preference for the lottery in Problem 14 implies $\pi(.001)v(5{,}000) > v(5)$, hence $\pi(.001) > v(5)/v(5{,}000) > .001$, assuming the value function for gains is concave. The readiness to pay for insurance in Problem 14′ implies the same conclusion, assuming the value function for losses is convex.

It is important to distinguish overweighting, which refers to a property of decision weights, from the overestimation that is commonly found in the assessment of the probability of rare events. Note that the issue of overestimation does not arise in the present context, where the subject is assumed to adopt the stated value of p. In many real-life situations, overestimation and overweighting may both operate to increase the impact of rare events.

Although $\pi(p) > p$ for low probabilities, there is evidence to suggest that, for all $0 < p < 1$, $\pi(p) + \pi(1-p) < 1$. We label this property subcertainty. It is readily seen that the typical preferences in any version of Allias' example (see, e.g., Problems 1 and 2) imply subcertainty for the relevant value of p. Applying equation (1) to the prevalent preferences in Problems 1 and 2 yields, respectively,

$$v(2{,}400) > \pi(.66)v(2{,}400) + \pi(.33)v(2{,}500), \quad \text{i.e.,}$$
$$[1 - \pi(.66)]v(2{,}400) > \pi(.33)v(2{,}500) \quad \text{and}$$
$$\pi(.33)v(2{,}500) > \pi(.34)v(2{,}400); \quad \text{hence,}$$
$$1 - \pi(.66) > \pi(.34), \quad \text{or} \quad \pi(.66) + \pi(.34) < 1.$$

Applying the same analysis to Allais' original example yields $\pi(.89) + \pi(.11) < 1$, and some data reported by MacCrimmon and Larsson (forthcoming) imply subcertainty for additional values of p.

The slope of π in the interval (0, 1) can be viewed as a measure of the sensitivity of preferences to changes in probability. Subcertainty entails that π is regressive with respect to p, i.e., that preferences are generally less sensitive to variations of probability than the expectation principle would dictate. Thus, subcertainty captures an essential element of people's attitudes to uncertain events, namely that the sum of the weights associated with complementary events is typically less than the weight associated with the certain event.

Recall that the violations of the substitution axiom discussed earlier in this paper conform to the following rule: If (x, p) is equivalent to (y, pq) then (x, pr) is not preferred to (y, pqr), $0 < p, q, r \leq 1$. By equation (1),

$$\pi(p)v(x) = \pi(pq)v(y) \quad \text{implies} \quad \pi(pr)v(x) \leq \pi(pqr)v(y); \quad \text{hence,}$$
$$\frac{\pi(pq)}{\pi(p)} \leq \frac{\pi(pqr)}{\pi(pr)}.$$

Thus, for a fixed ratio of probabilities, the ratio of the corresponding decision weights is closer to unity when the probabilities are low than when they are high.

This property of π, called subproportionality, imposes considerable constraints on the shape of π: it holds if and only if $\log \pi$ is a convex function of $\log p$.

It is of interest to note that subproportionality together with the overweighting of small probabilities imply that π is subadditive over that range. Formally, it can be shown that if $\pi(p) > p$ and subproportionality holds, then $\pi(rp) > r\pi(p)$, $0 < r < 1$, provided π is monotone and continuous over (0, 1).

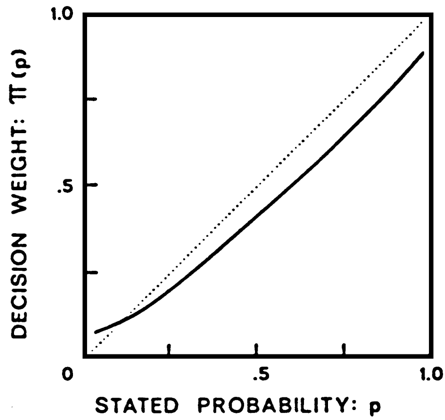

Figure 2.4. A hypothetical weighting function.

Figure 2.4 presents a hypothetical weighting function which satisfies overweighting and subadditivity for small values of p, as well as subcertainty and subproportionality. These properties entail that π is relatively shallow in the open interval and changes abruptly near the end-points where $\pi(0) = 0$ and $\pi(1) = 1$. The sharp drops or apparent discontinuities of π at the endpoints are consistent with the notion that there is a limit to how small a decision weight can be attached to an event, if it is given any weight at all. A similar quantum of doubt could impose an upper limit on any decision weight that is less than unity. This quantal effect may reflect the categorical distinction between certainty and uncertainty. On the other hand, the simplification of prospects in the editing phase can lead the individual to discard events of extremely low probability and to treat events of extremely high probability as if they were certain. Because people are limited in their ability to comprehend and evaluate extreme probabilities, highly unlikely events are either ignored or overweighted, and the difference between high probability and certainty is either neglected or exaggerated. Consequently, π is not well-behaved near the end-points.

The following example, due to Zeckhauser, illustrates the hypothesized nonlinearity of π. Suppose you are compelled to play Russian roulette, but are given the opportunity to purchase the removal of one bullet from the loaded gun. Would you pay as much to reduce the number of bullets from four to three as you would to reduce the number of bullets from one to zero? Most people feel that they would be willing to pay much more for a reduction of the probability of death from 1/6 to zero than for a reduction from 4/6 to 3/6. Economic considerations would lead one to pay more in the latter case, where the value of money is presumably reduced by the considerable probability that one will not live to enjoy it.

An obvious objection to the assumption that $\pi(p) \neq p$ involves comparisons between prospects of the form $(x, p; x, q)$ and $(x, p'; x, q')$, where $p + q = p' + q' < 1$. Since any individual will surely be indifferent between the two prospects, it could be argued that this observation entails $\pi(p) +$

$\pi(q) = \pi(p') + \pi(q')$, which in turn implies that π is the identity function. This argument is invalid in the present theory, which assumes that the probabilities of identical outcomes are combined in the editing of prospects. A more serious objection to the nonlinearity of π involves potential violations of dominance. Suppose $x > y > 0$, $p > p'$, and $p + q = p' + q' < 1$; hence, $(x, p; y, q)$ dominates $(x, p'; y, q')$. If preference obeys dominance, then

$$\pi(p)v(x) + \pi(q)v(y) > \pi(p')v(x) + \pi(q')v(y).$$

or

$$\frac{\pi(p) - \pi(p')}{\pi(q') - \pi(q)} > \frac{v(y)}{v(x)}.$$

Hence, as y approaches x, $\pi(p) - \pi(p')$ approaches $\pi(q') - \pi(q)$. Since $p - p' = q' - q$, π must be essentially linear, or else dominance must be violated.

Direct violations of dominance are prevented, in the present theory, by the assumption that dominated alternatives are detected and eliminated prior to the evaluation of prospects. However, the theory permits indirect violations of dominance, e.g., triples of prospects so that A is preferred to B, B is preferred to C, and C dominates A. For an example, see Raiffa [1968, p. 75].

Finally, it should be noted that the present treatment concerns the simplest decision task in which a person chooses between two available prospects. We have not treated in detail the more complicated production task (e.g., bidding) where the decision maker generates an alternative that is equal in value to a given prospect. The asymmetry between the two options in this situation could introduce systematic biases. Indeed, Lichtenstein and Slovic [1971] have constructed pairs of prospect A and B, such that people generally prefer A over B, but bid more for B than A. This phenomenon has been confirmed in several studies, with both hypothetical and real gambles, e.g., Grether and Plott (forthcoming). Thus, it cannot be generally assumed that the preference order of prospects can be recovered by a bidding procedure.

Because prospect theory has been proposed as a model of choice, the inconsistency of bids and choices implies that the measurement of values and decision weights should be based on choices between specified prospects rather than on bids or other production tasks. This restriction makes the assessment of v and π more difficult because production tasks are more convenient for scaling than pair comparisons.

4. DISCUSSION

In the final section we show how prospect theory accounts for observed attitudes toward risk, discuss alternative representations of choice problems induced by shifts of reference point, and sketch several extensions of the present treatment.

Risk Attitudes

The dominant pattern of preferences observed in Allais' example (Problems 1 and 2) follows from the present theory iff

$$\frac{\pi(.33)}{\pi(.34)} > \frac{v(2{,}400)}{v(2{,}500)} > \frac{\pi(.33)}{1-\pi(.66)}.$$

Hence, the violation of the independence axiom is attributed in this case to subcertainty, and more specifically to the inequality $\pi(.34) < 1 - \pi(.66)$. This analysis shows that an Allais-type violation will occur whenever the v-ratio of the two non-zero outcomes is bounded by the corresponding π-ratios.

Problems 3 through 8 share the same structure, hence it suffices to consider one pair, say Problems 7 and 8. The observed choices in these problems are implied by the theory iff

$$\frac{\pi(.001)}{\pi(.002)} > \frac{v(3{,}000)}{v(6{,}000)} > \frac{\pi(.45)}{\pi(.90)}.$$

The violation of the substitution axiom is attributed in this case to the subproportionality of π. Expected utility theory is violated in the above manner, therefore, whenever the v-ratio of the two outcomes is bounded by the respective π-ratios. The same analysis applies to other violations of the substitution axiom, both in the positive and in the negative domain.

We next prove that the preference for regular insurance over probabilistic insurance, observed in Problem 9, follows from prospect theory – provided the probability of loss is overweighted. That is, if $(-x, p)$ is indifferent to $(-y)$, then $(-y)$ is preferred to $(-x, p/2; -y, p/2; -y/2, 1-p)$. For simplicity, we define for $x \geq 0$, $f(x) = -v(-x)$. Since the value function for losses is convex, f is a concave function of x. Applying prospect theory, with the natural extension of equation 2, we wish to show that

$$\begin{aligned} &\pi(p)f(x) = f(y) \quad \text{implies} \\ &f(y) \leq f(y/2) + \pi(p/2)[f(y) - f(y/2)] + \pi(p/2)[f(x) - f(y/2)] \\ &\qquad = \pi(p/2)f(x) + \pi(p/2)f(y) + [1 - 2\pi(p/2)]f(y/2). \end{aligned}$$

Substituting for $f(x)$ and using the concavity of f, it suffices to show that

$$f(y) \leq \frac{\pi(p/2)}{\pi(p)} f(y) + \pi(p/2)f(y) + f(y)/2 - \pi(p/2)f(y)$$

or

$$\pi(p)/2 \leq \pi(p/2), \quad \text{which follows from the subadditivity of } \pi.$$

According to the present theory, attitudes toward risk are determined jointly by v and π, and not solely by the utility function. It is therefore instructive to examine the conditions under which risk aversion or risk seeking are expected to occur. Consider the choice between the gamble (x, p) and its expected value (px). If $x > 0$, risk seeking is implied whenever $\pi(p) > v(px)/v(x)$, which is greater than p if the value function for gains is concave. Hence, overweighting ($\pi(p) > p$) is necessary but not sufficient for risk seeking in the domain of gains. Precisely the same condition is necessary but not sufficient for risk aversion when $x < 0$. This analysis restricts risk seeking in the domain of gains and risk aversion in the domain of losses to small probabilities, where overweighting is expected to hold. Indeed these are the typical conditions under which lottery tickets and insurance policies are sold. In prospect theory, the overweighting of small probabilities favors both gambling and insurance, while the *S*-shaped value function tends to inhibit both behaviors.

Although prospect theory predicts both insurance and gambling for small probabilities, we feel that the present analysis falls far short of a fully adequate account of these complex phenomena. Indeed, there is evidence from both experimental studies (Slovic, Fischhoff, Lichtenstein, Corrigan, & Coombs 1977), survey research (Kunreuther et al., 1978), and observations of economic behavior, e.g., service and medical insurance, that the purchase of insurance often extends to the medium range of probabilities, and that small probabilities of disaster are sometimes entirely ignored. Furthermore, the evidence suggests that minor changes in the formulation of the decision problem can have marked effects on the attractiveness of insurance (Slovic, Fischhoff, Lichtenstein, Corrigan and Coombs, 1977). A comprehensive theory of insurance behavior should consider, in addition to pure attitudes toward uncertainty and money, such factors as the value of security, social norms of prudence, the aversiveness of a large number of small payments spread over time, information and misinformation regarding probabilities and outcomes, and many others. Some effects of these variables could be described within the present framework, e.g., as changes of reference point, transformations of the value function, or manipulations of probabilities or decision weights. Other effects may require the introduction of variables or concepts which have not been considered in this treatment.

Shifts of Reference

So far in this paper, gains and losses were defined by the amounts of money that are obtained or paid when a prospect is played, and the reference point was taken to be the status quo, or one's current assets. Although this is probably true for most choice problems, there are situations in which gains and losses are coded relative to an expectation or aspiration level that differs from the status quo. For example, an unexpected tax withdrawal from a monthly pay check is experienced as a loss, not as a reduced gain. Similarly, an entrepreneur who is weathering a slump with greater success than his competitors may interpret a small loss as a gain, relative to the larger loss he had reason to expect.

The reference point in the preceding examples corresponded to an asset position that one had expected to attain. A discrepancy between the reference point and the current asset position may also arise because of recent changes in wealth to which one has not yet adapted (Markowitz, 1952). Imagine a person who is involved in a business venture, has already lost 2,000 and is now facing a choice between a sure gain of 1,000 and an even chance to win 2,000 or nothing. If he has not yet adapted to his losses, he is likely to code the problem as a choice between $(-2,000, .50)$ and $(-1,000)$ rather than as a choice between (2,000, .50) and (1,000). As we have seen, the former representation induces more adventurous choices than the latter.

A change of reference point alters the preference order for prospects. In particular, the present theory implies that a negative translation of a choice problem, such as arises from incomplete adaptation to recent losses, increases risk seeking in some situations. Specifically, if a risky prospect $(x, p; -y, 1-p)$ is just acceptable, then $(x-z, p; -y-z, 1-p)$ is preferred over $(-z)$ for $x, y, z > 0$, with $x > z$.

To prove this proposition, note that

$$V(x, p; y, 1-p = 0 \quad \text{iff} \quad \pi(p)v(x) = -\pi(1-p)v(-y).$$

Furthermore,

$$\begin{aligned}
&V(x-z, p; -y-z, 1-p)\\
&\quad = \pi(p)v(x-z) + \pi(1-p)v(-y-z)\\
&\quad > \pi(p)v(x) - \pi(p)v(z) + \pi(1-p)v(-y)\\
&\qquad + \pi(1-p)v(-z) \quad \text{by the properties of } v,\\
&\quad = -\pi(1-p)v(-y) - \pi(p)v(z) + \pi(1-p)v(-y)\\
&\qquad + \pi(1-p)v(-z) \quad \text{by substitution,}\\
&\quad = -\pi(p)v(z) + \pi(1-p)v(-z)\\
&\quad > v(-z)[\pi(p) + \pi(1-p)]] \quad \text{since } v(-z) < -v(z),\\
&\quad > v(-z) \quad \text{by subcertainty.}
\end{aligned}$$

This analysis suggests that a person who has not made peace with his losses is likely to accept gambles that would be unacceptable to him otherwise. The well known observation (McGlothlin, 1956) that the tendency to bet on long shots increases in the course of the betting day provides some support for the hypothesis that a failure to adapt to losses or to attain an expected gain induces risk seeking. For another example, consider an individual who expects to purchase insurance, perhaps because he has owned it in the past or because his friends do. This individual may code the decision to pay a premium y to protect against a loss x as a choice between $(-x+y, p; y, 1-p)$ and (0) rather than as a choice between $(-x, p)$ and $(-y)$. The preceding argument entails that insurance is likely to be more attractive in the former representation than in the latter.

Another important case of a shift of reference point arises when a person formulates his decision problem in terms of final assets, as advocated in decision analysis, rather than in terms of gains and losses, as people usually do. In this case, the reference point is set to zero on the scale of wealth and the value function is likely to be concave everywhere (Spetzler, 1968). According to the present analysis, this formulation essentially eliminates risk seeking, except for gambling with low probabilities. The explicit formulation of decision problems in terms of final assets is perhaps the most effective procedure for eliminating risk seeking in the domain of losses.

Many economic decisions involve transactions in which one pays money in exchange for a desirable prospect. Current decision theories analyze such problems as comparisons between the status quo and an alternative state which includes the acquired prospect minus its cost. For example, the decision whether to pay 10 for the gamble (1,000, .01) is treated as a choice between (990, .01; –10, .99) and (0). In this analysis, readiness to purchase the positive prospect is equated to willingness to accept the corresponding mixed prospect.

The prevalent failure to integrate riskless and risky prospects, dramatized in the isolation effect, suggests that people are unlikely to perform the operation of subtracting the cost from the outcomes in deciding whether to buy a gamble. Instead, we suggest that people usually evaluate the gamble and its cost separately, and decide to purchase the gamble if the combined value is positive. Thus, the gamble (1,000, .01) will be purchased for a price of 10 if $\pi(.01)v(1{,}000) + v(-10) > 0$.

If this hypothesis is correct, the decision to pay 10 for (1,000, .01), for example, is no longer equivalent to the decision to accept the gamble (990, .01; –10, .99). Furthermore, prospect theory implies that if one is indifferent between $(x(1-p), p; -px, 1-p)$ and (0) then one will not pay px to purchase the prospect (x, p). Thus, people are expected to exhibit more risk seeking in deciding whether to accept a fair gamble than in deciding whether to purchase a gamble for a fair price. The location of the reference point, and the manner in which choice problems are coded and edited emerge as critical factors in the analysis of decisions.

Extensions

In order to encompass a wider range of decision problems, prospect theory should be extended in several directions. Some generalizations are immediate; others require further development. The extension of equations (1) and (2) to prospects with any number of outcomes is straightforward. When the number of outcomes is large, however, additional editing operations may be invoked to simplify evaluation. The manner in which complex options, e.g., compound prospects, are reduced to simpler ones is yet to be investigated.

Although the present paper has been concerned mainly with monetary outcomes, the theory is readily applicable to choices involving other attributes, e.g., quality of life or the number of lives that could be lost or saved as a consequence of a policy decision. The main properties of the proposed value function for money should apply to other attributes as well. In particular, we expect outcomes to be coded as gains or losses relative to a neutral reference point, and losses to loom larger than gains.

The theory can also be extended to the typical situation of choice, where the probabilities of outcomes are not explicitly given. In such situations, decision weights must be attached to particular events rather than to stated probabilities, but they are expected to exhibit the essential properties that were ascribed to the weighting function. For example, if A and B are complementary events and neither is certain, $\pi(A) + \pi(B)$ should be less than unity – a natural analogue to subcertainty.

The decision weight associated with an event will depend primarily on the perceived likelihood of that event, which could be subject to major biases (Tversky & Kahneman, 1974). In addition, decision weights may be affected by other considerations, such as ambiguity or vagueness. Indeed, the work of Ellsberg (1961) and Fellner (1965) implies that vagueness reduces decision weights. Consequently, subcertainty should be more pronounced for vague than for clear probabilities.

The present analysis of preference between risky options has developed two themes. The first theme concerns editing operations that determine how prospects are perceived. The second theme involves the judgmental principles that govern the evaluation of gains and losses and the weighting of uncertain outcomes. Although both themes should be developed further, they appear to provide a useful framework for the descriptive analysis of choice under risk.

3. Advances in Prospect Theory
Cumulative Representation of Uncertainty

Amos Tversky and Daniel Kahneman

ABSTRACT. We develop a new version of prospect theory that employs cumulative rather than separable decision weights and extends the theory in several respects. This version, called cumulative prospect theory, applies to uncertain as well as to risky prospects with any number of outcomes, and it allows different weighting functions for gains and for losses. Two principles, diminishing sensitivity and loss aversion, are invoked to explain the characteristic curvature of the value function and the weighting functions. A review of the experimental evidence and the results of a new experiment confirm a distinctive fourfold pattern of risk attitudes: risk aversion for gains and risk seeking for losses of high probability; risk seeking for gains and risk aversion for losses of low probability.

KEY WORDS cumulative prospect theory

Expected utility theory reigned for several decades as the dominant normative and descriptive model of decision making under uncertainty, but it has come under serious question in recent years. There is now general agreement that the theory does not provide an adequate description of individual choice: a substantial body of evidence shows that decision makers systematically violate its basic tenets. Many alternative models have been proposed in response to this empirical challenge (for reviews, see Camerer, 1989; Fishburn, 1988; Machina, 1987). Some time ago we presented a model of choice, called prospect theory, which explained the major violations of expected utility theory in choices between risky prospects with a small number of outcomes (Kahneman and Tversky, 1979;

An earlier version of this article was entitled "Cumulative Prospect Theory: An Analysis of Decision under Uncertainty."

This article has benefited from discussions with Colin Camerer, Chew Soo-Hong, David Freedman, and David H. Krantz. We are especially grateful to Peter P. Wakker for his invaluable input and contribution to the axiomatic analysis. We are indebted to Richard Gonzalez and Amy Hayes for running the experiment and analyzing the data. This work was supported by Grants 89-0064 and 88-0206 from the Air Force Office of Scientific Research, by Grant SES-9109535 from the National Science Foundation, and by the Sloan Foundation.

Tversky and Kahneman, 1986). The key elements of this theory are 1) a value function that is concave for gains, convex for losses, and steeper for losses than for gains, and 2) a nonlinear transformation of the probability scale, which overweights small probabilities and underweights moderate and high probabilities. In an important later development, several authors (Quiggin, 1982; Schmeidler, 1989; Yaari, 1987; Weymark, 1981) have advanced a new representation, called the rank-dependent or the cumulative functional, that transforms cumulative rather than individual probabilities. This article presents a new version of prospect theory that incorporates the cumulative functional and extends the theory to uncertain as well to risky prospects with any number of outcomes. The resulting model, called cumulative prospect theory, combines some of the attractive features of both developments (see also Luce and Fishburn, 1991). It gives rise to different evaluations of gains and losses, which are not distinguished in the standard cumulative model, and it provides a unified treatment of both risk and uncertainty.

To set the stage for the present development, we first list five major phenomena of choice, which violate the standard model and set a minimal challenge that must be met by any adequate descriptive theory of choice. All these findings have been confirmed in a number of experiments, with both real and hypothetical payoffs.

Framing effects. The rational theory of choice assumes description invariance: equivalent formulations of a choice problem should give rise to the same preference order (Arrow, 1982). Contrary to this assumption, there is much evidence that variations in the framing of options (e.g., in terms of gains or losses) yield systematically different preferences (Tversky and Kahneman, 1986).

Nonlinear preferences. According to the expectation principle, the utility of a risky prospect is linear in outcome probabilities. Allais's (1953) famous example challenged this principle by showing that the difference between probabilities of .99 and 1.00 has more impact on preferences than the difference between 0.10 and 0.11. More recent studies observed nonlinear preferences in choices that do not involve sure things (Camerer and Ho, 1991).

Source dependence. People's willingness to bet on an uncertain event depends not only on the degree of uncertainty but also on its source. Ellsberg (1961) observed that people prefer to bet on an urn containing equal numbers of red and green balls, rather than on an urn that contains red and green balls in unknown proportions. More recent evidence indicates that people often prefer a bet on an event in their area of competence over a bet on a matched chance event, although the former probability is vague and the latter is clear (Heath and Tversky, 1991).

Risk seeking. Risk aversion is generally assumed in economic analyses of decision under uncertainty. However, risk-seeking choices are consistently observed in two classes of decision problems. First, people often prefer a small probability of winning a large prize over the expected value of that prospect. Second, risk seeking is prevalent when people must choose between a sure loss and a substantial probability of a larger loss.

Loss aversion. One of the basic phenomena of choice under both risk and uncertainty is that losses loom larger than gains (Kahneman and Tversky, 1984; Tversky and Kahneman, 1991). The observed asymmetry between gains and losses is far too extreme to be explained by income effects or by decreasing risk aversion.

The present development explains loss aversion, risk seeking, and nonlinear preferences in terms of the value and the weighting functions. It incorporates a framing process, and it can accommodate source preferences. Additional phenomena that lie beyond the scope of the theory – and of its alternatives – are discussed later.

The present article is organized as follows. Section 1.1 introduces the (two-part) cumulative functional; section 1.2 discusses relations to previous work; and section 1.3 describes the qualitative properties of the value and the weighting functions. These properties are tested in an extensive study of individual choice, described in section 2, which also addresses the question of monetary incentives. Implications and limitations of the theory are discussed in section 3.

1. THEORY

Prospect theory distinguishes two phases in the choice process: framing and valuation. In the framing phase, the decision maker constructs a representation of the acts, contingencies, and outcomes that are relevant to the decision. In the valuation phase, the decision maker assesses the value of each prospect and chooses accordingly. Although no formal theory of framing is available, we have learned a fair amount about the rules that govern the representation of acts, outcomes, and contingencies (Tversky and Kahneman, 1986). The valuation process discussed in subsequent sections is applied to framed prospects.

1.1. Cumulative Prospect Theory

In the classical theory, the utility of an uncertain prospect is the sum of the utilities of the outcomes, each weighted by its probability. The empirical evidence reviewed above suggests two major modifications of this theory: 1) the carriers of value are gains and losses, not final assets; and 2) the value of each outcome is multiplied by a decision weight, not by an additive probability. The weighting scheme used in the original version of prospect theory and in other models is a monotonic transformation of outcome probabilities. This scheme encounters two problems. First, it does not always satisfy stochastic dominance, an assumption that many theorists are reluctant to give up. Second, it is not readily extended to prospects with a large number of outcomes. These problems can be handled by assuming that transparently dominated prospects are eliminated in the editing phase, and by normalizing the weights so that they add to unity. Alternatively, both problems can be solved by the rank-dependent or cumulative functional, first proposed by Quiggin (1982) for decision under risk and by Schmeidler (1989) for decision under uncertainty. Instead of transforming each probability separately, this model transforms the entire cumulative distribution function. The present theory applies the cumulative functional separately to gains and to losses. This development extends prospect theory to uncertain as well as to risky prospects with any number of outcomes while preserving most of its essential features. The differences between the cumulative and the original versions of the theory are discussed in section 1.2.

Let S be a finite set of states of nature; subsets of S are called events. It is assumed that exactly one state obtains, which is unknown to the decision maker. Let X be a set of consequences, also called outcomes. For simplicity, we confine the present dicussion to monetary outcomes. We assume that X includes a neutral outcome, denoted 0, and we interpret all other elements of X as gains or losses, denoted by positive or negative numbers, respectively.

An uncertain prospect f is a function from S into X that assigns to each state $s \in S$ a consequence $f(s) = x$ in X. To define the cumulative functional, we arrange the outcomes of each prospect in increasing order. A prospect f is then represented as a sequence of pairs (x_i, A_i), which yields x_i if A_i occurs, where $x_i > x_j$ iff $i > j$, and (A_i) is a partition of S. We use positive subscripts to denote positive outcomes, negative subscripts to denote negative outcomes, and the zero subscript to index the neutral outcome. A prospect is called strictly positive or positive, respectively, if its outcomes are all positive or nonnegative. Strictly negative and negative prospects are defined similarly; all other prospects are called mixed. The positive part of f, denoted f^+, is obtained by letting $f^+(s) = f(s)$ if $f(s) > 0$, and $f^+(s) = 0$ if $f(s) \leq 0$. The negative part of f, denoted f^-, is defined similarly.

As in expected utility theory, we assign to each prospect f a number $V(f)$ such that f is preferred to or indifferent to g iff $V(f) \geq V(g)$. The following representation is defined in terms of the concept of *capacity* (Choquet, 1955), a nonadditive set function that generalizes the standard notion of probability. A capacity W is a function that assigns to each $A \subset S$ a number $W(A)$ satisfying $W(\phi) = 0$, $W(S) = 1$, and $W(A) \geq W(B)$ whenever $A \supset B$.

Cumulative prospect theory asserts that there exist a strictly increasing value function $v : X \to \mathrm{Re}$, satisfying $v(x_0) = v(0) = 0$, and capacities W^+ and W^-, such that for $f = (x_i, A_i)$, $-m \leq i \leq n$,

$$V(f) = V(f^+) + V(f^-),$$

$$V(f^+) = \sum_{i=0}^{n} \pi_i^+ v(x_i), \quad V(f^-) = \sum_{i=-m}^{0} \pi_i^- v(x_i), \tag{1}$$

where the decision weights $\pi^+(f^+) = (\pi_0^+, \ldots, \pi_n^+)$ and $\pi^-(f^-) = (\pi_{-m}^-, \ldots, \pi_0^-)$ are defined by:

$$\begin{aligned}
\pi_n^+ &= W^+(A_n), \quad \pi_{-m}^- = W^-(A_{-m}),\\
\pi_i^+ &= W^+(A_i \cup \cdots \cup A_n) - W^+(A_{i+1} \cup \cdots \cup A_n), \quad 0 \le i \le n-1,\\
\pi_i^- &= W^-(A_{-m} \cup \cdots \cup A_i) - W^-(A_{-m} \cup \cdots \cup A_{i-1}), \quad 1-m \le i \le 0.
\end{aligned}$$

Letting $\pi_i = \pi_i^+$ if $i \ge 0$ and $\pi_i = \pi_i^-$ if $i < 0$, equation (1) reduces to

$$V(f) = \sum_{i=-m}^{n} \pi_i v(x_i). \tag{2}$$

The decision weight π_i^+, associated with a positive outcome, is the difference between the capacities of the events "the outcome is at least as good as x_i" and "the outcome is strictly better than x_i." The decision weight π_i^-, associated with a negative outcome, is the difference between the capacities of the events "the outcome is at least as bad as x_i" and "the outcome is strictly worse than x_i." Thus, the decision weight associated with an outcome can be interpreted as the marginal contribution of the respective event,[1] defined in terms of the capacities W^+ and W^-. If each W is additive, and hence a probability measure, then π_i is simply the probability of A_i. It follows readily from the definitions of π and W that for both positive and negative prospects, the decision weights add to 1. For mixed prospects, however, the sum can be either smaller or greater than 1, because the decision weights for gains and for losses are defined by separate capacities.

If the prospect $f = (x_i, A_i)$ is given by a probability distribution $p(A_i) = p_i$, it can be viewed as a probabilistic or risky prospect (x_i, p_i). In this case, decision weights are defined by:

$$\begin{aligned}
\pi_n^+ &= w^+(p_n), \quad \pi_{-m}^- = w^-(p_{-m}),\\
\pi_i^+ &= w^+(p_i + \cdots + p_n) - w^+(p_{i+1} + \cdots + p_n), \quad 0 \le i \le n-1,\\
\pi_i^- &= w^-(p_{-m} + \cdots + p_i) - w^-(p_{-m} + \cdots + p_{i-1}), \quad 1-m \le i \le 0,
\end{aligned}$$

where w^+ and w^- are strictly increasing functions from the unit interval into itself satisfying $w^+(0) = w^-(0) = 0$, and $w^+(1) = w^-(1) = 1$.

To illustrate the model, consider the following game of chance. You roll a die once and observe the result $x = 1, \ldots, 6$. If x is even, you receive \$$x$; if x is odd, you pay \$$x$. Viewed as a probabilistic prospect with equiprobable outcomes, f yields the consequences $(-5, -3, -1, 2, 4, 6)$, each with probability $1/6$.

[1] In keeping with the spirit of prospect theory, we use the decumulative form for gains and the cumulative form for losses. This notation is vindicated by the experimental findings described in section 2.

Thus, $f^+ = (0, 1/2; 2, 1/6; 4, 1/6; 6, 1/6)$, and $f^- = (-5, 1/6; -3, 1/6; -1, 1/6; 0, 1/2)$. By equation (1), therefore,

$$\begin{aligned} V(f) &= V(f^+) + V(f^-) \\ &= v(2)[w^+(1/2) - w^+(1/3)] + v(4)[w^+(1/3) - w^+(1/6)] \\ &\quad + v(6)[w^+(1/6) - w^+(0)] + v(-5)[w^-(1/6) - w^-(0)] \\ &\quad + v(-3)[w^-(1/3) - w^-(1/6)] + v(-1)[w^-(1/2) - w^-(1/3)]. \end{aligned}$$

1.2. Relation to Previous Work

Luce and Fishburn (1991) derived essentially the same representation from a more elaborate theory involving an operation $\circ$ of joint receipt or multiple play. Thus, $f \circ g$ is the composite prospect obtained by playing both f and g, separately. The key feature of their theory is that the utility function U is additive with respect to $\circ$, that is, $U(f \circ g) = U(f) + U(g)$ provided one prospect is acceptable (i.e., preferred to the status quo) and the other is not. This condition seems too restrictive both normatively and descriptively. As noted by the authors, it implies that the utility of money is a linear function of money if for all sums of money x, y, $U(x \circ y) = U(x + y)$. This assumption appears to us inescapable because the joint receipt of x and y is tantamount to receiving their sum. Thus, we expect the decision maker to be indifferent between receiving a \$10 bill or receiving a \$20 bill and returning \$10 in change. The Luce–Fishburn theory, therefore, differs from ours in two essential respects. First, it extends to composite prospects that are not treated in the present theory. Second, it practically forces utility to be proportional to money.

The present representation encompasses several previous theories that employ the same decision weights for all outcomes. Starmer and Sugden (1989) considered a model in which $w^-(p) = w^+(p)$, as in the original version of prospect theory. In contrast, the rank-dependent models assume $w^-(p) = 1 - w^+(1 - p)$ or $W^-(A) = 1 - W^+(S - A)$. If we apply the latter condition to choice between uncertain assets, we obtain the choice model established by Schmeidler (1989), which is based on the Choquet integral.[2] Other axiomatizations of this model were developed by Gilboa (1987), Nakamura (1990), and Wakker (1989a, 1989b). For probabilistic (rather than uncertain) prospects, this model was first established by Quiggin (1982) and Yaari (1987), and was further analyzed by Chew (1989), Segal (1989), and Wakker (1990). An earlier axiomatization of this model in the context of income inequality was presented by Weymark (1981). Note that in the present theory, the overall value $V(f)$ of a mixed prospect is not a Choquet integral but rather a sum $V(f^+) + V(f^-)$ of two such integrals.

[2] This model appears under different names. We use *cumulative utility theory* to describe the application of a Choquet integral to a standard utility function, and *cumulative prospect theory* to describe the application of two separate Choquet integrals to the value of gains and losses.

The present treatment extends the original version of prospect theory in several respects. First, it applies to any finite prospect and it can be extended to continuous distributions. Second, it applies to both probabilistic and uncertain prospects and can, therefore, accommodate some form of source dependence. Third, the present theory allows different decision weights for gains and losses, thereby generalizing the original version that assumes $w^+ = w^-$. Under this assumption, the present theory coincides with the original version for all two-outcome prospects and for all mixed three-outcome prospects. It is noteworthy that for prospects of the form $(x, p; y, 1-p)$, where either $x > y > 0$ or $x < y < 0$, the original theory is in fact rank dependent. Although the two models yield similar predictions in general, the cumulative version – unlike the original one – satisfies stochastic dominance. Thus, it is no longer necessary to assume that transparently dominated prospects are eliminated in the editing phase – an assumption that was criticized by some authors. On the other hand, the present version can no longer explain violations of stochastic dominance in nontransparent contexts (e.g., Tversky and Kahneman, 1986).

1.3. Values and Weights

In expected utility theory, risk aversion and risk seeking are determined solely by the utility function. In the present theory, as in other cumulative models, risk aversion and risk seeking are determined jointly by the value function and by the capacities, which in the present context are called cumulative weighting functions, or weighting functions for short. As in the original version of prospect theory, we assume that v is concave above the reference point ($v''(x) \leq 0$, $x \geq 0$) and convex below the reference point ($v''(x) \geq 0$, $x \leq 0$). We also assume that v is steeper for losses than for gains $v'(x) < v'(-x)$ for $x \geq 0$. The first two conditions reflect the principle of diminishing sensitivity: the impact of a change diminishes with the distance from the reference point. The last condition is implied by the principle of loss aversion according to which losses loom larger than corresponding gains (Tversky and Kahneman, 1991).

The principle of diminishing sensitivity applies to the weighting functions as well. In the evaluation of outcomes, the reference point serves as a boundary that distinguishes gains from losses. In the evaluation of uncertainty, there are two natural boundaries – certainty and impossibility – that correspond to the endpoints of the certainty scale. Diminishing sensitivity entails that the impact of a given change in probability diminishes with its distance from the boundary. For example, an increase of .1 in the probability of winning a given prize has more impact when it changes the probability of winning from .9 to 1.0 or from 0 to .1, than when it changes the probability of winning from .3 to .4 or from .6 to .7. Diminishing sensitivity, therefore, gives rise to a weighting function that is concave near 0 and convex near 1. For uncertain prospects, this principle yields subadditivity for very unlikely events and superadditivity near certainty. However, the function is not well-behaved near the endpoints, and very small probabilities can be either greatly overweighted or neglected altogether.

Table 3.1. A Test of Independence (Dow-Jones)

		A	B	C	
		if $d < 30$	**if $30 \leq d \leq 35$**	**if $35 < d$**	
Problem I:	f	$25,000	$25,000	$25,000	[68]
	g	$25,000	0	$75,000	[32]
Problem II:	f'	0	$25,000	$25,000	[23]
	g'	0	0	$75,000	[77]

Note: Outcomes are contingent on the difference d between the closing values of the Dow-Jones today and tomorrow. The percentage of respondents ($N = 156$) who selected each prospect is given in brackets.

Before we turn to the main experiment, we wish to relate the observed nonlinearity of preferences to the shape of the weighting function. For this purpose, we devised a new demonstration of the common consequence effect in decisions involving uncertainty rather than risk. Table 3.1 displays a pair of decision problems (I and II) presented in that order to a group of 156 money managers during a workshop. The participants chose between prospects whose outcomes were contingent on the difference d between the closing values of the Dow-Jones today and tomorrow. For example, f' pays $25,000 if d exceeds 30 and nothing otherwise. The percentage of respondents who chose each prospect is given in brackets. The independence axiom of expected utility theory implies that f is preferred to g iff f' is preferred to g'. Table 3.1 shows that the modal choice was f in problem I and g' in problem II. This pattern, which violates independence, was chosen by 53% of the respondents.

Essentially the same pattern was observed in a second study following the same design. A group of 98 Stanford students chose between prospects whose outcomes were contingent on the point-spread d in the forthcoming Stanford–Berkeley football game. Table 3.2 presents the prospects in question. For example, g pays $10 if Stanford does not win, $30 if it wins by 10 points or less, and nothing if it wins by more than 10 points. Ten percent of the participants, selected at random, were actually paid according to one of their choices. The

Table 3.2. A Test of Independence (Stanford-Berkeley football game)

		A	B	C	
		if $d < 0$	**if $0 \leq d \leq 10$**	**if $10 < d$**	
Problem I:	f	$10	$10	$10	[64]
	g	$10	$30	0	[36]
Problem II:	f'	0	$10	$10	[34]
	g'	0	$30	0	[66]

Note: Outcomes are contingent on the point-spread d in a Stanford–Berkeley football game. The percentage of respondents ($N = 98$) who selected each prospect is given in brackets.

modal choice, selected by 46% of the subjects, was f and g', again in direct violation of the independence axiom.

To explore the constraints imposed by this pattern, let us apply the present theory to the modal choices in Table 3.1, using $1,000 as a unit. Since f is preferred to g in problem I,

$$v(25) > v(75)W^+(C) + v(25)[W^+(A \cup C) - W^+(C)]$$

or

$$v(25)[1 - W^+(A \cup C) + W^+(C)] > v(75)W^+(C).$$

The preference for g' over f' in problem II, however, implies

$$v(75)W^+(C) > v(25)W^+(C \cup B);$$

hence,

$$W^+(S) - W^+(S - B) > W^+(C \cup B) - W^+(C). \quad (3)$$

Thus, "subtracting" B from certainty has more impact than "subtracting" B from $C \cup B$. Let $W_+(D) = 1 - W^+(S - D)$, and $w_+(p) = 1 - w^+(1 - p)$. It follows readily that equation (3) is equivalent to the subadditivity of W_+, that is, $W_+(B) + W_+(D) \geq W_+(B \cup D)$. For probabilistic prospects, equation (3) reduces to

$$1 - w^+(1 - q) > w^+(p + q) - w^+(p),$$

or

$$w_+(q) + w_+(r) \geq w_+(q + r), \quad q + r < 1.$$

Allais's example corresponds to the case where $p(C) = .10$, $p(B) = .89$, and $p(A) = .01$.

It is noteworthy that the violations of independence reported in Tables 3.1 and 3.2 are also inconsistent with regret theory, advanced by Loomes and Sugden (1982, 1987), and with Fishburn's (1988) SSA model. Regret theory explains Allais's example by assuming that the decision maker evaluates the consequences as if the two prospects in each choice are statistically independent. When the prospects in question are defined by the same set of events, as in Tables 3.1 and 3.2, regret theory (like Fishburn's SSA model) implies independence, since it is additive over states. The finding that the common consequence effect is very much in evidence in the present problems undermines the interpretation of Allais's example in terms of regret theory.

The common consequence effect implies the subadditivity of W_+ and of w_+. Other violations of expected utility theory imply the subadditivity of W^+ and of w^+ for small and moderate probabilities. For example, Prelec (1990) observed that most respondents prefer 2% to win $20,000 over 1% to win $30,000; they also prefer 1% to win $30,000 and 32% to win $20,000 over

34% to win $20,000. In terms of the present theory, these data imply that $w^+(.02) - w^+(.01) \geq w^+(.34) - w^+(.33)$. More generally, we hypothesize

$$w^+(p+q) - w^+(q) \geq w^+(p+q+r) - w^+(q+r), \tag{4}$$

provided $p+q+r$ is sufficiently small. Equation (4) states that w^+ is concave near the origin; and the conjunction of the above inequalities implies that, in accord with diminishing sensitivity, w^+ has an inverted S-shape: it is steepest near the endpoints and shallower in the middle of the range. For other treatments of decision weights, see Hogarth and Einhorn (1990), Prelec (1989), Viscusi (1989), and Wakker (1990). Experimental evidence is presented in the next section.

2. EXPERIMENT

An experiment was carried out to obtain detailed information about the value and weighting functions. We made a special effort to obtain high-quality data. To this end, we recruited 25 graduate students from Berkeley and Stanford (12 men and 13 women) with no special training in decision theory. Each subject participated in three separate one-hour sessions that were several days apart. Each subject was paid $25 for participation.

2.1. Procedure

The experiment was conducted on a computer. On a typical trial, the computer displayed a prospect (e.g., 25% chance to win $150 and 75% chance to win $50) and its expected value. The display also included a descending series of seven sure outcomes (gains or losses) logarithmically spaced between the extreme outcomes of the prospect. The subject indicated a preference between each of the seven sure outcomes and the risky prospect. To obtain a more refined estimate of the certainty equivalent, a new set of seven sure outcomes was then shown, linearly spaced between a value 25% higher than the lowest amount accepted in the first set and a value 25% lower than the highest amount rejected. The certainty equivalent of a prospect was estimated by the midpoint between the lowest accepted value and the highest rejected value in the second set of choices. We wish to emphasize that although the analysis is based on certainty equivalents, the data consisted of a series of choices between a given prospect and several sure outcomes. Thus, the cash equivalent of a prospect was derived from observed choices, rather than assessed by the subject. The computer monitored the internal consistency of the responses to each prospect and rejected errors, such as the acceptance of a cash amount lower than one previously rejected. Errors caused the original statement of the problem to reappear on the screen.

The present analysis focuses on a set of two-outcome prospects with monetary outcomes and numerical probabilities. Other data involving more complicated prospects, including prospects defined by uncertain events, will be reported elsewhere. There were 28 positive and 28 negative prospects. Six of the prospects (three nonnegative and three nonpositive) were repeated on

Table 3.3. Median Cash Equivalents (in dollars) for All Nonmixed Prospects

	Probability								
Outcomes	**.01**	**.05**	**.10**	**.25**	**.50**	**.75**	**.90**	**.95**	**.99**
(0, 50)			9		21		37		
(0, −50)			−8		−21		−39		
(0, 100)		14		25	36	52		78	
(0, −100)		−8		−23.5	−42	−63		−84	
(0, 200)	10		20		76		131		188
(0, −200)	−3		−23		−89		−155		−190
(0, 400)	12								377
(0, −400)	−14								−380
(50, 100)			59		71		83		
(−50, −100)			−59		−71		−85		
(50, 150)		64		72.5	86	102		128	
(−50, −150)		−60		−71	−92	−113		−132	
(100, 200)		118		130	141	162		178	
(−100, −200)		−112		−121	−142	−158		−179	

Note: The two outcomes of each prospect are given in the left-hand side of each row; the probability of the second (i.e., more extreme) outcome is given by the corresponding column. For example, the value of $9 in the upper left corner is the median cash equivalent of the prospect (0, 9; $50, .1).

different sessions to obtain the estimate of the consistency of choice. Table 3.3 displays the prospects and the median cash equivalents of the 25 subjects.

A modified procedure was used in eight additional problems. In four of these problems, the subjects made choices regarding the acceptability of a set of mixed prospects (e.g., 50% chance to lose $100 and 50% chance to win x) in which x was systematically varied. In four other problems, the subjects compared a fixed prospect (e.g., 50% chance to lose $20 and 50% chance to win $50) to a set of prospects (e.g., 50% chance to lose $50 and 50% chance to win x) in which x was systematically varied. (These prospects are presented in Table 3.6.)

2.2. Results

The most distinctive implication of prospect theory is the fourfold pattern of risk attitudes. For the nonmixed prospects used in the present study, the shapes of the value and the weighting functions imply risk-averse and risk-seeking preferences, respectively, for gains and for losses of moderate or high probability. Furthermore, the shape of the weighting functions favors risk seeking for small probabilities of gains and risk aversion for small probabilities of loss, provided the outcomes are not extreme. Note, however, that prospect theory does not imply perfect reflection in the sense that the preference between any two positive prospects is reversed when gains are replaced by losses. Table 3.4 presents, for each subject, the percentage of risk-seeking choices (where the certainty equivalent exceeded expected value) for gains and for losses with low

Table 3.4. Percentage of Risk-Seeking Choices

	Gain		Loss	
Subject	$p \leq .1$	$p \geq .5$	$p \leq .1$	$p \geq .5$
1	100	38	30	100
2	85	33	20	75
3	100	10	0	93
4	71	0	30	58
5	83	0	20	100
6	100	5	0	100
7	100	10	30	86
8	87	0	10	100
9	16	0	80	100
10	83	0	0	93
11	100	26	0	100
12	100	16	10	100
13	87	0	10	94
14	100	21	30	100
15	66	0	30	100
16	60	5	10	100
17	100	15	20	100
18	100	22	10	93
19	60	10	60	63
20	100	5	0	81
21	100	0	0	100
22	100	0	0	92
23	100	31	0	100
24	71	0	80	100
25	100	0	10	87
Risk seeking	78[a]	10	20	87[a]
Risk neutral	12	2	0	7
Risk averse	10	88[a]	80[a]	6

[a] Values that correspond to the fourfold pattern.
Note: The percentage of risk-seeking choices is given for low ($p \leq .1$) and high ($p \geq .5$) probabilities of gain and loss for each subject (risk-neutral choices were excluded). The overall percentage of risk-seeking, risk-neutral, and risk-averse choices for each type of prospect appears at the bottom of the table.

($p \leq .1$) and with high ($p \geq .5$) probabilities. Table 3.4 shows that for $p \geq .5$, all 25 subjects are predominantly risk averse for positive prospects and risk seeking for negative ones. Moreover, the entire fourfold pattern is observed for 22 of the 25 subjects, with some variability at the level of individual choices.

Although the overall pattern of preferences is clear, the individual data, of course, reveal both noise and individual differences. The correlations, across subjects, between the cash equivalents for the same prospects on successive sessions averaged .55 over six different prospects. Table 3.5 presents means

Table 3.5. Average Correlations between Certainty Equivalents in Four Types of Prospects

	L^+	**H^+**	**L^-**	**H^-**
L^+	.41	.17	−.23	.05
H^+		.39	.05	−.18
L^-			.40	.06
H^-				.44

Note: Low probability of gain = L^+; high probability of gain = H^+; low probability of loss = L^-; high probability of loss = H^-.

(after transformation to Fisher's z) of the correlations between the different types of prospects. For example, there were 19 and 17 prospects, respectively, with high probability of gain and high probability of loss. The value of .06 in Table 3.5 is the mean of the $17 \times 19 = 323$ correlations between the cash equivalents of these prospects.

The correlations between responses within each of the four types of prospects average .41, slightly lower than the correlations between separate responses to the same problems. The two negative values in Table 3.5 indicate that those subjects who were more risk averse in one domain tended to be more risk seeking in the other. Although the individual correlations are fairly low, the trend is consistent: 78% of the 403 correlations in these two cells are negative. There is also a tendency for subjects who are more risk averse for high-probability gains to be less risk seeking for gains of low probability. This trend, which is absent in the negative domain, could reflect individual differences either in the elevation of the weighting function or in the curvature of the value function for gains. The very low correlations in the two remaining cells of Table 3.5, averaging .05, indicate that there is no general trait of risk aversion or risk seeking. Because individual choices are quite noisy, aggregation of problems is necessary for the analysis of individual differences.

The fourfold pattern of risk attitudes emerges as a major empirical generalization about choice under risk. It has been observed in several experiments (see, e.g., Cohen, Jaffray, and Said, 1987), including a study of experienced oil executives involving significant, albeit hypothetical, gains and losses (Wehrung, 1989). It should be noted that prospect theory implies the pattern demonstrated in Table 3.4 within the data of individual subjects, but it does not imply high correlations across subjects because the values of gains and of losses can vary independently. The failure to appreciate this point and the limited reliability of individual responses has led some previous authors (e.g., Hershey and Schoemaker, 1980) to underestimate the robustness of the fourfold pattern.

2.3. Scaling

Having established the fourfold pattern in ordinal and correlational analyses, we now turn to a quantitative description of the data. For each prospect of the form $(x, p; 0, 1-p)$, let c/x be the ratio of the certainty equivalent of the prospect to the nonzero outcome x. Figures 3.1 and 3.2 plot the median value of c/x as a function of p, for positive and for negative prospects, respectively. We denote c/x by a circle if $|x| < 200$, and by a triangle if $|x| \geq 200$. The only exceptions are the two extreme probabilities (.01 and .99) where a circle is used for $|x| = 200$. To

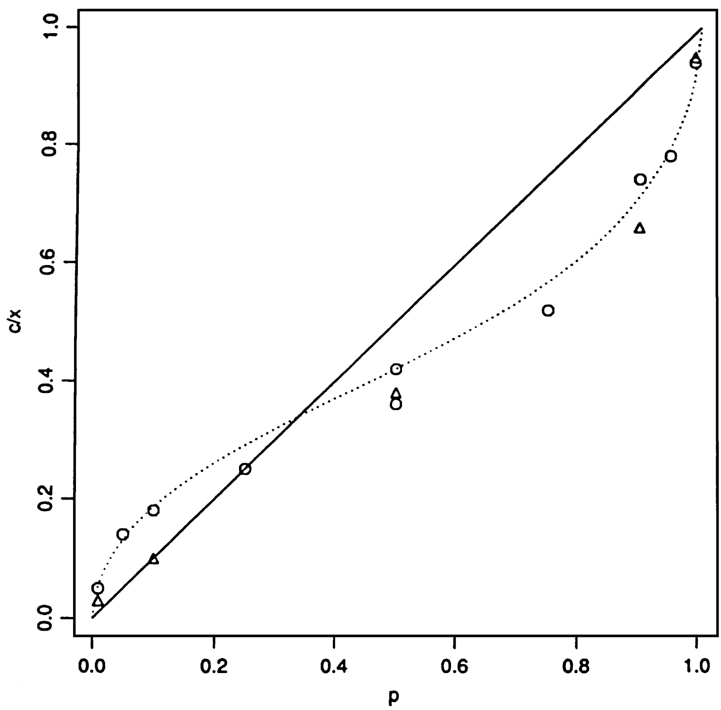

Figure 3.1. Median c/x for all positive prospects of the form $(x, p; 0, 1-p)$. Triangles and circles, respectively, correspond to values of x that lie above or below 200.

interpret Figures 3.1 and 3.2, note that if subjects are risk neutral, the points will lie on the diagonal; if subjects are risk averse, all points will lie below the diagonal in Figure 3.1 and above the diagonal in Figure 3.2. Finally, the triangles and the circles will lie on top of each other if preferences are homogeneous, so that multiplying the outcomes of a prospect f by a constant $k > 0$ multiplies its cash equivalent $c(kf)$ by the same constant, that is, $c(kf) = kc(f)$. In expected utility theory, preference homogeneity gives rise to constant relative risk aversion. Under the present theory, assuming $X = \text{Re}$, preference homogeneity is both necessary and sufficient to represent v as a two-part power function of the form

$$v(x) = \begin{cases} x^{\alpha} & \text{if } x \geq 0 \\ -\lambda(-x)^{\beta} & \text{if } x < 0. \end{cases} \tag{5}$$

Figures 3.1 and 3.2 exhibit the characteristic pattern of risk aversion and risk seeking observed in Table 3.4. They also indicate that preference homogeneity holds as a good approximation. The slight departures from homogeneity in Figure 3.1 suggest that the cash equivalents of positive prospects increase more slowly than the stakes (triangles tend to lie below the circles), but no such tendency is evident in Figure 3.2. Overall, it appears that the present data can be

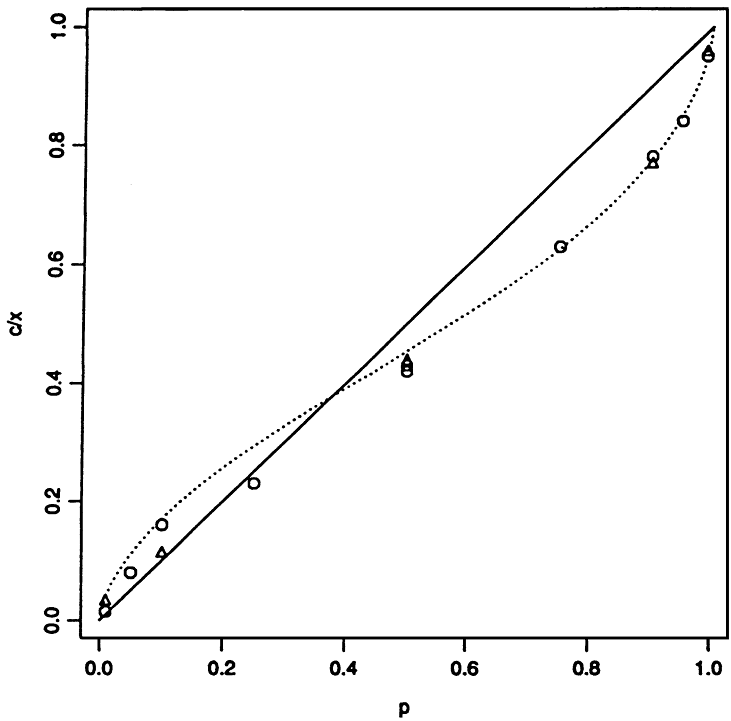

Figure 3.2. Median c/x for all negative prospects of the form $(x, p; 0, 1-p$. Triangles and circles, respectively, correspond to values of x that lie below or above -200.

approximated by a two-part power function. The smooth curves in Figures 3.1 and 3.2 can be interpreted as weighting functions, assuming a linear value function. They were fitted using the following functional form:

$$w^+(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{1/\gamma}}, \qquad w^-(p) = \frac{p^\delta}{(p^\delta + (1-p)^\delta)^{1/\delta}}. \tag{6}$$

This form has several useful features: it has only one parameter; it encompasses weighting functions with both concave and convex regions; it does not require $w(.5) = .5$; and most important, it provides a reasonably good approximation to both the aggregate and the individual data for probabilities in the range between .05 and .95.

Further information about the properties of the value function can be derived from the data presented in Table 3.6. The adjustments of mixed prospects to acceptability (problems 1–4) indicate that, for even chances to win and lose, a prospect will only be acceptable if the gain is at least twice as large as the loss. This observation is compatible with a value function that changes slope abruptly, at zero, with a loss-aversion coefficient of about 2 (Tversky and Kahneman, 1991). The median matches in problems 5 and 6 are also consistent with this estimate: when the possible loss is increased by k the compensating

Table 3.6. A Test of Loss Aversion

Problem	a	b	c	x	θ
1	0	0	−25	61	2.44
2	0	0	−50	101	2.02
3	0	0	−100	202	2.02
4	0	0	−150	280	1.87
5	−20	50	−50	112	2.07
6	−50	150	−125	301	2.01
7	50	120	20	149	0.97
8	100	300	25	401	1.35

Note: In each problem, subjects determined the value of x that makes the prospect ($\$a$, $^1/_2$; $\$b$, $^1/_2$) as attractive as ($\$c$, $^1/_2$; $\$x$, $^1/_2$). The median values of x are presented for all problems along with the fixed values a, b, c. The statistic $\theta = (x - b)/(c - a)$ is the ratio of the "slopes" at a higher and a lower region of the value function.

gain must be increased by about $2k$. Problems 7 and 8 are obtained from problems 5 and 6, respectively, by positive translations that turn mixed prospects into strictly positive ones. In contrast to the large values of θ observed in problems 1–6, the responses in problems 7 and 8 indicate that the curvature of the value function for gains is slight. A decrease in the smallest gain of a strictly positive prospect is fully compensated by a slightly larger increase in the largest gain. The standard rank-dependent model, which lacks the notion of a reference point, cannot account for the dramatic effects of small translations of prospects illustrated in Table 3.6.

The estimation of a complex choice model, such as cumulative prospect theory, is problematic. If the functions associated with the theory are not constrained, the number of estimated parameters for each subject is too large. To reduce this number, it is common to assume a parametric form (e.g., a power utility function), but this approach confounds the general test of the theory with that of the specific parametric form. For this reason, we focused here on the qualitative properties of the data rather than on parameter estimates and measures of fit. However, in order to obtain a parsimonious description of the present data, we used a nonlinear regression procedure to estimate the parameters of equations (5) and (6), separately for each subject. The median exponent of the value function was 0.88 for both gains and losses, in accord with diminishing sensitivity. The median λ was 2.25, indicating pronounced loss aversion, and the median values of γ and δ, respectively, were 0.61 and 0.69, in agreement with equations (3) and (4) above.[3] The parameters estimated from the median data were essentially the same. Figure 3 plots w^+ and w^- using the median estimates of γ and δ.

[3] Camerer and Ho (1991) applied equation (6) to several studies of risky choice and estimated γ from aggregate choice probabilities using a logistic distribution function. Their mean estimate (.56) was quite close to ours.

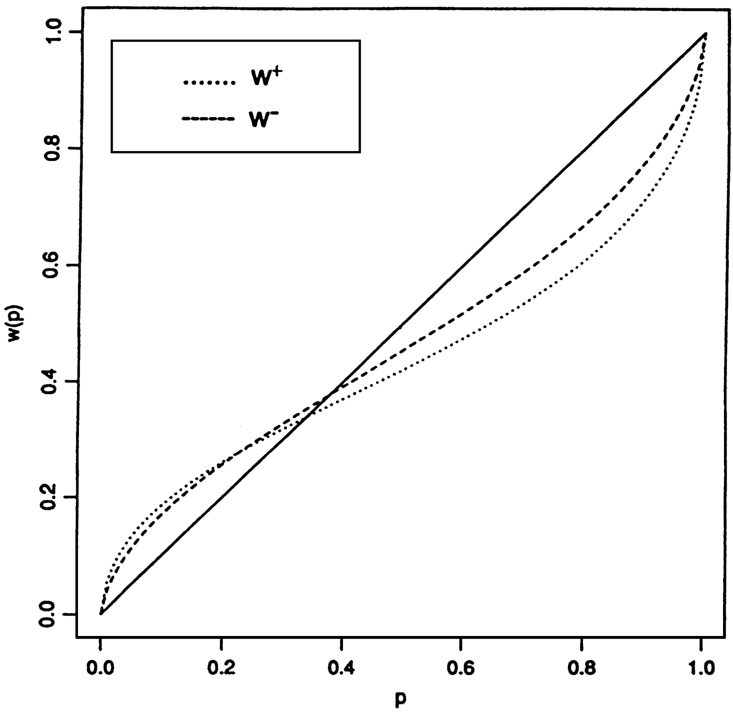

Figure 3.3. Weighting functions for gains (w^+) and for losses (w^-) based on median estimates of γ and δ in equation (12).

Figure 3.3 shows that, for both positive and negative prospects, people overweight low probabilities and underweight moderate and high probabilities. As a consequence, people are relatively insensitive to probability difference in the middle of the range. Figure 3.3 also shows that the weighting functions for gains and for losses are quite close, although the former is slightly more curved than the latter (i.e., $\gamma < \delta$). Accordingly, risk aversion for gains is more pronounced than risk seeking for losses, for moderate and high probabilities (see Table 3.3). It is noteworthy that the condition $w^+(p) = w^-(p)$, assumed in the original version of prospect theory, accounts for the present data better than the assumption $w^+(p) = 1 - w^-(1-p)$, implied by the standard rank-dependent or cumulative functional. For example, our estimates of w^+ and w^- show that all 25 subjects satisfied the conditions $w^+(.5) < .5$ and $w^-(.5) < .5$, implied by the former model, and no one satisfied the condition $w^+(.5) < .5$ iff $w^-(.5) > .5$, implied by the latter model.

Much research on choice between risky prospects has utilized the triangle diagram (Marschak, 1950; Machina, 1987) that represents the set of all prospects of the form $(x_1, p_1; x_2, p_2; x_3, p_3)$, with fixed outcomes $x_1 < x_2 < x_3$. Each point in the triangle represents a prospect that yields the lowest outcome (x_1) with probability p_1, the highest outcome (x_3) with probability p_3, and the intermediate outcome (x_2) with probability $p_2 = 1 - p_1 - p_3$. An indifference curve is a set of prospects (i.e., points) that the decision maker finds equally attractive. Alternative choice theories are characterized by the shapes of their indifference curves. In particular, the indifference curves of expected utility theory are parallel straight lines. Figures 3.4a and 3.4b illustrate the indifference curves of cumulative prospect theory for nonnegative and nonpositive prospects, respectively. The shapes of the curves are determined by the weighting functions of Figure 3.3; the values of the outcomes (x_1, x_2, x_3) merely control the slope.

Figures 3.4a and 3.4b are in general agreement with the main empirical generalizations that have emerged from the studies of the triangle diagram; see Camerer (1992), and Camerer and Ho (1991) for reviews. First, departures from linearity, which violate expected utility theory, are most pronounced near the edges of the triangle. Second, the indifference curves exhibit both fanning in and fanning out. Third, the curves are concave in the upper part of the triangle and convex in the lower right. Finally, the indifference curves for nonpositive prospects resemble the curves for nonnegative prospects reflected around the 45° line, which represents risk neutrality. For example, a sure gain of \$100 is equally as attractive as a 71% chance to win \$200 or nothing (see Figure 3.4a), and a sure loss of \$100 is equally as aversive as a 64% chance to lose \$200 or nothing (see Figure 3.4b). The approximate reflection of the curves is of special interest because it distinguishes the present theory from the standard rank-dependent model in which the two sets of curves are essentially the same.

2.4. Incentives

We conclude this section with a brief discussion of the role of monetary incentives. In the present study we did not pay subjects on the basis of their choices because in our experience with choice between prospects of the type used in the present study, we did not find much difference between subjects who were paid a flat fee and subjects whose payoffs were contingent on their decisions. The same conclusion was obtained by Camerer (1989), who investigated the effects of incentives using several hundred subjects. He found that subjects who actually played the gamble gave essentially the same responses as subjects who did not play; he also found no differences in reliability and roughly the same decision time. Although some studies found differences between paid and unpaid subjects in choice between simple prospects, these differences were not large enough to change any significant qualitative conclusions. Indeed, all major violations of expected utility theory (e.g. the common consequence effect, the common ratio effect, source dependence, loss aversion, and preference reversals) were obtained both with and without monetary incentives.

As noted by several authors, however, the financial incentives provided in choice experiments are generally small relative to people's incomes. What

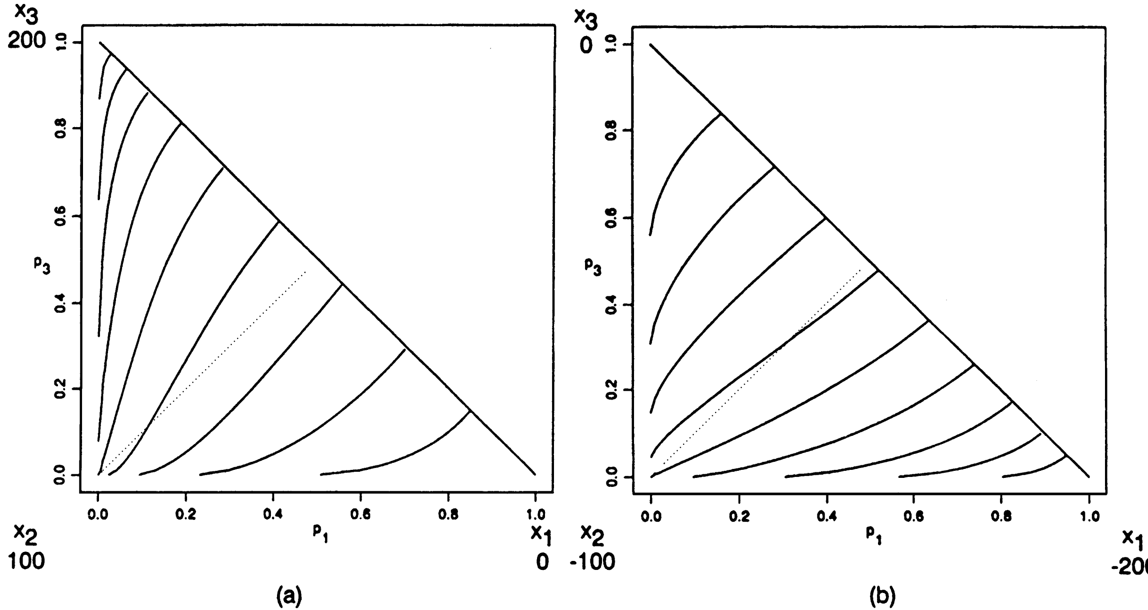

Figure 3.4. Indifference curves of cumulative prospect theory (a) for nonnegative prospects ($x_1 = 0$, $x_2 = 100$, $x_3 = 200$), and (b) for nonpositive prospects ($x_1 = -200, x_2 = -100$, $x_3 = 0$). The curves are based on the respective weighting functions of Figure 3.3, ($\gamma = .61$, $\delta = .69$) and on the median estimates of the exponents of the value function ($\alpha = \beta = .88$). The broken line through the origin represents the prospects whose expected value is x_2.

happens when the stakes correspond to three- or four-digit rather than one- or two-digit figures? To answer this question, Kachelmeier and Shehata (1991) conducted a series of experiments using Masters students at Beijing University, most of whom had taken at least one course in economics or business. Due to the economic conditions in China, the investigators were able to offer subjects very large rewards. In the high payoff condition, subjects earned about three times their normal monthly income in the course of one experimental session! On each trial, subjects were presented with a simple bet that offered a specified probability to win a given prize, and nothing otherwise. Subjects were instructed to state their cash equivalent for each bet. An incentive-compatible procedure (the BDM scheme) was used to determine, on each trial, whether the subject would play the bet or receive the "official" selling price. If departures from the standard theory are due to the mental cost associated with decision making and the absence of proper incentives, as suggested by Smith and Walker (1992), then the highly paid Chinese subjects should not exhibit the characteristic nonlinearity observed in hypothetical choices, or in choices with small payoffs.

However, the main finding of Kachelmeier and Shehata (1991) is massive risk seeking for small probabilities. Risk seeking was slightly more pronounced for lower payoffs, but even in the highest payoff condition, the cash equivalent for a 5% bet (their lowest probability level) was, on average, three times larger than its expected value. Note that in the present study the median cash equivalent of a 5% chance to win $100 (see Table 3.3) was $14, almost three times the expected value of the bet. In general, the cash equivalents obtained by Kachelmeier and Shehata were higher than those observed in the present study. This is consistent with the finding that minimal selling prices are generally higher than certainty equivalents derived from choice (see, e.g., Tversky, Slovic, and Kahneman, 1990). As a consequence, they found little risk aversion for moderate and high probability of winning. This was true for the Chinese subjects, at both high and low payoffs, as well as for Canadian subjects, who either played for low stakes or did not receive any payoff. The most striking result in all groups was the marked overweighting of small probabilities, in accord with the present analysis.

Evidently, high incentives do not always dominate noneconomic considerations, and the observed departures from expected utility theory cannot be rationalized in terms of the cost of thinking. We agree with Smith and Walker (1992) that monetary incentives could improve performance under certain conditions by eliminating careless errors. However, we maintain that monetary incentives are neither necessary nor sufficient to ensure subjects' cooperativeness, thoughtfulness, or truthfulness. The similarity between the results obtained with and without monetary incentives in choice between simple prospects provides no special reason for skepticism about experiments without contingent payment.

3. DISCUSSION

Theories of choice under uncertainty commonly specify 1) the objects of choice, 2) a valuation rule, and 3) the characteristics of the functions that map uncertain events and possible outcomes into their subjective counterparts. In standard applications of expected utility theory, the objects of choice are probability distributions over wealth, the valuation rule is expected utility, and utility is a concave function of wealth. The empirical evidence reported here and elsewhere requires major revisions of all three elements. We have proposed an alternative descriptive theory in which 1) the objects of choice are prospects framed in terms of gains and losses, 2) the valuation rule is a two-part cumulative functional, and 3) the value function is S-shaped and the weighting functions are inverse S-shaped. The experimental findings confirmed the qualitative properties of these scales, which can be approximated by a (two-part) power value function and by identical weighting functions for gains and losses.

The curvature of the weighting function explains the characteristic reflection pattern of attitudes to risky prospects. Overweighting of small probabilities contributes to the popularity of both lotteries and insurance. Underweighting of high probabilities contributes both to the prevalence of risk aversion in choices between probable gains and sure things, and to the prevalence of risk seeking in choices between probable and sure losses. Risk aversion for gains and risk seeking for losses are further enhanced by the curvature of the value function in the two domains. The pronounced asymmetry of the value function, which we have labeled loss aversion, explains the extreme reluctance to accept mixed prospects. The shape of the weighting function explains the certainty effect and violations of quasi-convexity. It also explains why these phenomena are most readily observed at the two ends of the probability scale, where the curvature of the weighting function is most pronounced (Camerer, 1992).

The new demonstrations of the common consequence effect, described in Tables 3.1 and 3.2, show that choice under uncertainty exhibits some of the main characteristics observed in choice under risk. On the other hand, there are indications that the decision weights associated with uncertain and with risky prospects differ in important ways. First, there is abundant evidence that subjective judgments of probability do not conform to the rules of probability theory (Kahneman, Slovic and Tversky, 1982). Second, Ellsberg's example and more recent studies of choice under uncertainty indicate that people prefer some sources of uncertainty over others. For example, Heath and Tversky (1991) found that individuals consistently preferred bets on uncertain events in their area of expertise over matched bets on chance devices, although the former are ambiguous and the latter are not. The presence of systematic preferences for some sources of uncertainty calls for different weighting functions for different domains, and suggests that some of these functions lie entirely above others. The investigation of decision weights for uncertain events emerges as a promising domain for future research.

The present theory retains the major features of the original version of prospect theory and introduces a (two-part) cumulative functional, which provides a convenient mathematical representation of decision weights. It also relaxes some descriptively inappropriate constraints of expected utility theory. Despite its greater generality, the cumulative functional is unlikely to be accurate in detail. We suspect that decision weights may be sensitive to the formulation of the prospects, as well as to the number, the spacing and the level of outcomes. In particular, there is some evidence to suggest that the curvature of the weighting function is more pronounced when the outcomes are widely spaced (Camerer, 1992). The present theory can be generalized to accommodate such effects, but it is questionable whether the gain in descriptive validity, achieved by giving up the separability of values and weights, would justify the loss of predictive power and the cost of increased complexity.

Theories of choice are at best approximate and incomplete. One reason for this pessimistic assessment is that choice is a constructive and contingent process. When faced with a complex problem, people employ a variety of heuristic procedures in order to simplify the representation and the evaluation of prospects. These procedures include computational shortcuts and editing operations, such as eliminating common components and discarding nonessential differences (Tversky, 1969). The heuristics of choice do not readily lend themselves to formal analysis because their application depends on the formulation of the problem, the method of elicitation, and the context of choice.

Prospect theory departs from the tradition that assumes the rationality of economic agents; it is proposed as a descriptive, not a normative, theory. The idealized assumption of rationality in economic theory is commonly justified on two grounds: the conviction that only rational behavior can survive in a competitive environment, and the fear that any treatment that abandons rationality will be chaotic and intractable. Both arguments are questionable. First, the evidence indicates that people can spend a lifetime in a competitive environment without acquiring a general ability to avoid framing effects or to apply linear decision weights. Second, and perhaps more important, the evidence indicates that human choices are orderly, although not always rational in the traditional sense of this word.