Rational Choice and the Framing of Decisions
Author(s): Amos Tversky and Daniel Kahneman
Source: *The Journal of Business*, Vol. 59, No. 4, Part 2: The Behavioral Foundations of Economic Theory (Oct., 1986), pp. S251-S278
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/2352759
Accessed: 10/09/2008 09:33

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at http://www.jstor.org/page/info/about/policies/terms.jsp. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at http://www.jstor.org/action/showPublisher?publisherCode=ucpress.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit organization founded in 1995 to build trusted digital archives for scholarship. We work with the scholarly community to preserve their work and the materials they rely upon, and to build a common research platform that promotes the discovery and use of these resources. For more information about JSTOR, please contact support@jstor.org.

http://www.jstor.org

Amos Tversky
Stanford University

Daniel Kahneman
University of British Columbia

Rational Choice and the Framing of Decisions*

Alternative descriptions of a decision problem often give rise to different preferences, contrary to the principle of invariance that underlies the rational theory of choice. Violations of this theory are traced to the rules that govern the framing of decision and to the psychophysical principles of evaluation embodied in prospect theory. Invariance and dominance are obeyed when their application is transparent and often violated in other situations. Because these rules are normatively essential but descriptively invalid, no theory of choice can be both normatively adequate and descriptively accurate.

The modern theory of decision making under risk emerged from a logical analysis of games of chance rather than from a psychological analysis of risk and value. The theory was conceived as a normative model of an idealized decision maker, not as a description of the behavior of real people. In Schumpeter's words, it "has a much better claim to being called a logic of choice than a psychology of value" (1954, p. 1058).

The use of a normative analysis to predict and explain actual behavior is defended by several arguments. First, people are generally thought to be effective in pursuing their goals, particularly when they have incentives and opportunities to learn from experience. It seems reasonable, then, to describe choice as a maximization process. Second, competition favors rational individuals and organizations. Optimal decisions increase the chances of survival in a competitive environment, and a minority of rational individuals can sometimes impose rationality on the

* This work was supported by contract N00014-84-K-0615 from the Office of Naval Research to Stanford University. The present article reviews our work on decision making under risk from a new perspective, discussed primarily in the first and last sections. Most of the empirical demonstrations have been reported in earlier publications. Problems 3, 4, 7, 8, and 12 are published here for the first time. Requests for reprints should be addressed to Amos Tversky, Department of Psychology, Stanford University, Stanford, California 94705, or to Daniel Kahneman, Department of Psychology, University of California, Berkeley, California 94720.

(*Journal of Business*, 1986, vol. 59, no. 4, pt. 2)
© 1986 by The University of Chicago. All rights reserved.
0021-9398/86/5904-0010$01.50

whole market. Third, the intuitive appeal of the axioms of rational choice makes it plausible that the theory derived from these axioms should provide an acceptable account of choice behavior.

The thesis of the present article is that, in spite of these a priori arguments, the logic of choice does not provide an adequate foundation for a descriptive theory of decision making. We argue that the deviations of actual behavior from the normative model are too widespread to be ignored, too systematic to be dismissed as random error, and too fundamental to be accommodated by relaxing the normative system. We first sketch an analysis of the foundations of the theory of rational choice and then show that the most basic rules of the theory are commonly violated by decision makers. We conclude from these findings that the normative and the descriptive analyses cannot be reconciled. A descriptive model of choice is presented, which accounts for preferences that are anomalous in the normative theory.

I. A Hierarchy of Normative Rules

The major achievement of the modern theory of decision under risk is the derivation of the expected utility rule from simple principles of rational choice that make no reference to long-run considerations (von Neumann and Morgenstern 1944). The axiomatic analysis of the foundations of expected utility theory reveals four substantive assumptions—cancellation, transitivity, dominance, and invariance—besides the more technical assumptions of comparability and continuity. The substantive assumptions can be ordered by their normative appeal, from the cancellation condition, which has been challenged by many theorists, to invariance, which has been accepted by all. We briefly discuss these assumptions.

Cancellation. The key qualitative property that gives rise to expected utility theory is the "cancellation" or elimination of any state of the world that yields the same outcome regardless of one's choice. This notion has been captured by different formal properties, such as the substitution axiom of von Neumann and Morgenstern (1944), the extended sure-thing principle of Savage (1954), and the independence condition of Luce and Krantz (1971). Thus, if A is preferred to B, then the prospect of winning A if it rains tomorrow (and nothing otherwise) should be preferred to the prospect of winning B if it rains tomorrow because the two prospects yield the same outcome (nothing) if there is no rain tomorrow. Cancellation is necessary to represent preference between prospects as the maximization of expected utility. The main argument for cancellation is that only one state will actually be realized, which makes it reasonable to evaluate the outcomes of options separately for each state. The choice between options should therefore depend only on states in which they yield different outcomes.

Transitivity. A basic assumption in models of both risky and riskless choice is the transitivity of preference. This assumption is necessary and essentially sufficient for the representation of preference by an ordinal utility scale u such that A is preferred to B whenever $u(A) > u(B)$. Thus transitivity is satisfied if it is possible to assign to each option a value that does not depend on the other available options. Transitivity is likely to hold when the options are evaluated separately but not when the consequences of an option depend on the alternative to which it is compared, as implied, for example, by considerations of regret. A common argument for transitivity is that cyclic preferences can support a "money pump," in which the intransitive person is induced to pay for a series of exchanges that returns to the initial option.

Dominance. This is perhaps the most obvious principle of rational choice: if one option is better than another in one state and at least as good in all other states, the dominant option should be chosen. A slightly stronger condition—called stochastic dominance—asserts that, for unidimensional risky prospects, A is preferred to B if the cumulative distribution of A is to the right of the cumulative distribution of B. Dominance is both simpler and more compelling than cancellation and transitivity, and it serves as the cornerstone of the normative theory of choice.

Invariance. An essential condition for a theory of choice that claims normative status is the principle of invariance: different representations of the same choice problem should yield the same preference. That is, the preference between options should be independent of their description. Two characterizations that the decision maker, on reflection, would view as alternative descriptions of the same problem should lead to the same choice—even without the benefit of such reflection. This principle of invariance (or extensionality [Arrow 1982]), is so basic that it is tacitly assumed in the characterization of options rather than explicitly stated as a testable axiom. For example, decision models that describe the objects of choice as random variables all assume that alternative representations of the same random variables should be treated alike. Invariance captures the normative intuition that variations of form that do not affect the actual outcomes should not affect the choice. A related concept, called consequentialism, has been discussed by Hammond (1985).

The four principles underlying expected utility theory can be ordered by their normative appeal. Invariance and dominance seem essential, transitivity could be questioned, and cancellation has been rejected by many authors. Indeed, the ingenious counterexamples of Allais (1953) and Ellsberg (1961) led several theorists to abandon cancellation and the expectation principle in favor of more general representations. Most of these models assume transitivity, dominance, and invariance

(e.g., Hansson 1975; Allais 1979; Hagen 1979; Machina 1982; Quiggin 1982; Weber 1982; Chew 1983; Fishburn 1983; Schmeidler 1984; Segal 1984; Yaari 1984; Luce and Narens 1985). Other developments abandon transitivity but maintain invariance and dominance (e.g., Bell 1982; Fishburn 1982, 1984; Loomes and Sugden 1982). These theorists responded to observed violations of cancellation and transitivity by weakening the normative theory in order to retain its status as a descriptive model. However, this strategy cannot be extended to the failures of dominance and invariance that we shall document. Because invariance and dominance are normatively essential and descriptively invalid, a theory of rational decision cannot provide an adequate description of choice behavior.

We next illustrate failures of invariance and dominance and then review a descriptive analysis that traces these failures to the joint effects of the rules that govern the framing of prospects, the evaluation of outcomes, and the weighting of probabilities. Several phenomena of choice that support the present account are described.

II. Failures of Invariance

In this section we consider two illustrative examples in which the condition of invariance is violated and discuss some of the factors that produce these violations.

The first example comes from a study of preferences between medical treatments (McNeil et al. 1982). Respondents were given statistical information about the outcomes of two treatments of lung cancer. The same statistics were presented to some respondents in terms of mortality rates and to others in terms of survival rates. The respondents then indicated their preferred treatment. The information was presented as follows.[1]

Problem 1 (Survival frame)

Surgery: Of 100 people having surgery 90 live through the post-operative period, 68 are alive at the end of the first year and 34 are alive at the end of five years.

Radiation Therapy: Of 100 people having radiation therapy all live through the treatment, 77 are alive at the end of one year and 22 are alive at the end of five years.

Problem 1 (Mortality frame)

Surgery: Of 100 people having surgery 10 die during surgery or the post-operative period, 32 die by the end of the first year and 66 die by the end of five years.

1. All problems are presented in the text exactly as they were presented to the participants in the experiments.

Radiation Therapy: Of 100 people having radiation therapy, none die during treatment, 23 die by the end of one year and 78 die by the end of five years.

The inconsequential difference in formulation produced a marked effect. The overall percentage of respondents who favored radiation therapy rose from 18% in the survival frame ($N = 247$) to 44% in the mortality frame ($N = 336$). The advantage of radiation therapy over surgery evidently looms larger when stated as a reduction of the risk of immediate death from 10% to 0% rather than as an increase from 90% to 100% in the rate of survival. The framing effect was not smaller for experienced physicians or for statistically sophisticated business students than for a group of clinic patients.

Our next example concerns decisions between conjunctions of risky prospects with monetary outcomes. Each respondent made two choices, one between favorable prospects and one between unfavorable prospects (Tversky and Kahneman 1981, p. 454). It was assumed that the two selected prospects would be played independently.

Problem 2 ($N = 150$). 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%]

The total number of respondents is denoted by N, and the percentage who chose each option is indicated in brackets. (Unless otherwise specified, the data were obtained from undergraduate students at Stanford University and at the University of British Columbia.) The majority choice in decision i is risk averse, while the majority choice in decision ii is risk seeking. This is a common pattern: choices involving gains are usually risk averse, and choices involving losses are often risk seeking—except when the probability of winning or losing is small (Fishburn and Kochenberger 1979; Kahneman and Tversky 1979; Hershey and Schoemaker 1980).

Because the subjects considered the two decisions simultaneously, they expressed, in effect, a preference for the portfolio A and D over the portfolio B and C. However, the preferred portfolio is actually dominated by the rejected one! The combined options are as follows.

A & D: 25% chance to win $240 and 75% chance to lose $760.
B & C: 25% chance to win $250 and 75% chance to lose $750.

When the options are presented in this aggregated form, the dominant option is invariably chosen. In the format of problem 2, however, 73% of respondents chose the dominated combination A and D, and only 3% chose B and C. The contrast between the two formats illustrates a violation of invariance. The findings also support the general point that failures of invariance are likely to produce violations of stochastic dominance and vice versa.

The respondents evidently evaluated decisions i and ii separately in problem 2, where they exhibited the standard pattern of risk aversion in gains and risk seeking in losses. People who are given these problems are very surprised to learn that the combination of two preferences that they considered quite reasonable led them to select a dominated option. The same pattern of results was also observed in a scaled-down version of problem 2, with real monetary payoff (see Tversky and Kahneman 1981, p. 458).

As illustrated by the preceding examples, variations in the framing of decision problems produce systematic violations of invariance and dominance that cannot be defended on normative grounds. It is instructive to examine two mechanisms that could ensure the invariance of preferences: canonical representations and the use of expected actuarial value.

Invariance would hold if all formulations of the same prospect were transformed to a standard canonical representation (e.g., a cumulative probability distribution of the same random variable) because the various versions would then all be evaluated in the same manner. In problem 2, for example, invariance and dominance would both be preserved if the outcomes of the two decisions were aggregated prior to evaluation. Similarly, the same choice would be made in both versions of the medical problem if the outcomes were coded in terms of one dominant frame (e.g., rate of survival). The observed failures of invariance indicate that people do not spontaneously aggregate concurrent prospects or transform all outcomes into a common frame.

The failure to construct a canonical representation in decision problems contrasts with other cognitive tasks in which such representations are generated automatically and effortlessly. In particular, our visual experience consists largely of canonical representations: objects do not appear to change in size, shape, brightness, or color when we move around them or when illumination varies. A white circle seen from a sharp angle in dim light appears circular and white, not ellipsoid and grey. Canonical representations are also generated in the process of language comprehension, where listeners quickly recode much of what they hear into an abstract propositional form that no longer discriminates, for example, between the active and the passive voice and often does not distinguish what was actually said from what was implied or presupposed (Clark and Clark 1977). Unfortunately, the mental ma-

chinery that transforms percepts and sentences into standard forms does not automatically apply to the process of choice.

Invariance could be satisfied even in the absence of a canonical representation if the evaluation of prospects were separately linear, or nearly linear, in probability and monetary value. If people ordered risky prospects by their actuarial values, invariance and dominance would always hold. In particular, there would be no difference between the mortality and the survival versions of the medical problem. Because the evaluation of outcomes and probabilities is generally nonlinear, and because people do not spontaneously construct canonical representations of decisions, invariance commonly fails. Normative models of choice, which assume invariance, therefore cannot provide an adequate descriptive account of choice behavior. In the next section we present a descriptive account of risky choice, called prospect theory, and explore its consequences. Failures of invariance are explained by framing effects that control the representation of options, in conjunction with the nonlinearities of value and belief.

III. Framing and Evaluation of Outcomes

Prospect theory distinguishes two phases in the choice process: a phase of framing and editing, followed by a phase of evaluation (Kahneman and Tversky 1979). The first phase consists of a preliminary analysis of the decision problem, which frames the effective acts, contingencies, and outcomes. Framing is controlled by the manner in which the choice problem is presented as well as by norms, habits, and expectancies of the decision maker. Additional operations that are performed prior to evaluation include cancellation of common components and the elimination of options that are seen to be dominated by others. In the second phase, the framed prospects are evaluated, and the prospect of highest value is selected. The theory distinguishes two ways of choosing between prospects: by detecting that one dominates another or by comparing their values.

For simplicity, we confine the discussion to simple gambles with numerical probabilities and monetary outcomes. Let $(x, p; y, q)$ denote a prospect that yields x with probability p and y with probability q and that preserves the status quo with probability $(1 - p - q)$. According to prospect theory, there are values $v(\cdot)$, defined on gains and losses, and decision weights $\pi(\cdot)$, defined on stated probabilities, such that the overall value of the prospect equals $\pi(p)v(x) + \pi(q)v(y)$. A slight modification is required if all outcomes of a prospect have the same sign.[2]

2. If $p + q = 1$ and either $x > y > 0$ or $x < y < 0$, the value of a prospect is given by $v(y) + \pi(p)[v(x) - v(y)]$, so that decision weights are not applied to sure outcomes.

The Value Function

Following Markowitz (1952), outcomes are expressed in prospect theory as positive or negative deviations (gains or losses) from a neutral reference outcome, which is assigned a value of zero. Unlike Markowitz, however, we propose that the value function is commonly S shaped, concave above the reference point, and convex below it, as illustrated in figure 1. Thus the difference in subjective value between a gain of $100 and a gain of $200 is greater than the subjective difference between a gain of $1,100 and a gain of $1,200. The same relation between value differences holds for the corresponding losses. The proposed function expresses the property that the effect of a marginal change decreases with the distance from the reference point in either direction. These hypotheses regarding the typical shape of the value function may not apply to ruinous losses or to circumstances in which particular amounts assume special significance.

A significant property of the value function, called *loss aversion*, is that the response to losses is more extreme than the response to gains. The common reluctance to accept a fair bet on the toss of a coin suggests that the displeasure of losing a sum of money exceeds the pleasure of winning the same amount. Thus the proposed value function is (i) defined on gains and losses, (ii) generally concave for gains and convex for losses, and (iii) steeper for losses than for gains. These properties of the value function have been supported in many studies of risky choice involving monetary outcomes (Fishburn and Kochenberger 1979; Kahneman and Tversky 1979; Hershey and Schoemaker 1980; Payne, Laughhunn, and Crum 1980) and human lives (Tversky 1977; Eraker and Sox 1981; Tversky and Kahneman 1981; Fischhoff 1983). Loss aversion may also contribute to the observed discrepancies between the amount of money people are willing to pay for a good and the compensation they demand to give it up (Bishop and Heberlein 1979; Knetsch and Sinden 1984). This effect is implied by the value function if the good is valued as a gain in the former context and as a loss in the latter.

Framing Outcomes

The framing of outcomes and the contrast between traditional theory and the present analysis are illustrated in the following problems.

Problem 3 (N = 126): Assume yourself richer by $300 than you are today. You have to choose between
a sure gain of $100 [72%]
50% chance to gain $200 and 50% chance to gain nothing [28%]

Problem 4 (N = 128): Assume yourself richer by $500 than you are today. You have to choose between
a sure loss of $100 [36%]
50% chance to lose nothing and 50% chance to lose $200 [64%]

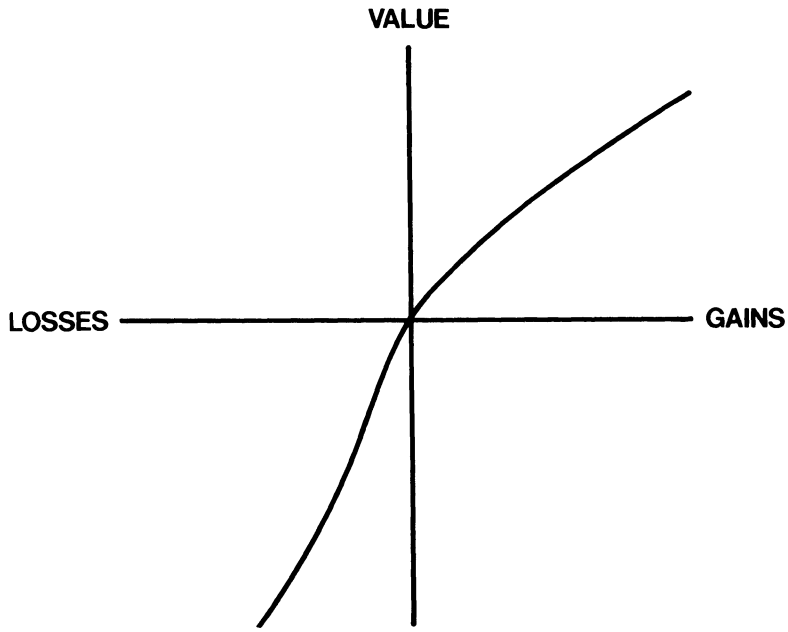

FIG. 1.—A typical value function

As implied by the value function, the majority choice is risk averse in problem 3 and risk seeking in problem 4, although the two problems are essentially identical. In both cases one faces a choice between $400 for sure and an even chance of $500 or $300. Problem 4 is obtained from problem 3 by increasing the initial endowment by $200 and subtracting this amount from both options. This variation has a substantial effect on preferences. Additional questions showed that variations of $200 in initial wealth have little or no effect on choices. Evidently, preferences are quite insensitive to small changes of wealth but highly sensitive to corresponding changes in reference point. These observations show that the effective carriers of values are gains and losses, or changes in wealth, rather than states of wealth as implied by the rational model.

The common pattern of preferences observed in problems 3 and 4 is of special interest because it violates not only expected utility theory but practically all other normatively based models of choice. In particular, these data are inconsistent with the model of regret advanced by Bell (1982) and by Loomes and Sugden (1982) and axiomatized by Fishburn (1982). This follows from the fact that problems 3 and 4 yield identical outcomes and an identical regret structure. Furthermore, regret theory cannot accommodate the combination of risk aversion in problem 3 and risk seeking in problem 4—even without the corresponding changes in endowment that make the problems extensionally equivalent.

Shifts of reference can be induced by different decompositions of outcomes into risky and riskless components, as in the above problems. The reference point can also be shifted by a mere labeling of outcomes, as illustrated in the following problems (Tversky and Kahneman 1981, p. 453).

Problem 5 (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 1/3 probability that 600 people will be saved, and 2/3 probability that no people will be saved. [28%]

In problem 5 the outcomes are stated in positive terms (lives saved), and the majority choice is accordingly risk averse. The prospect of certainly saving 200 lives is more attractive than a risky prospect of equal expected value. A second group of respondents was given the same cover story with the following descriptions of the alternative programs.

Problem 6 (N = 155):
If Program C is adopted 400 people will die. [22%]

If Program D is adopted there is 1/3 probability that nobody will die, and 2/3 probability that 600 people will die. [78%]

In problem 6 the outcomes are stated in negative terms (lives lost), and the majority choice is accordingly risk seeking. The certain death of 400 people is less acceptable than a two-thirds chance that 600 people will die. Problems 5 and 6, however, are essentially identical. They differ only in that the former is framed in terms of the number of lives saved (relative to an expected loss of 600 lives if no action is taken), whereas the latter is framed in terms of the number of lives lost.

On several occasions we presented both versions to the same respondents and discussed with them the inconsistent preferences evoked by the two frames. Many respondents expressed a wish to remain risk averse in the "lives saved" version and risk seeking in the "lives lost" version, although they also expressed a wish for their answers to be consistent. In the persistence of their appeal, framing effects resemble visual illusions more than computational errors.

Discounts and Surcharges

Perhaps the most distinctive intellectual contribution of economic analysis is the systematic consideration of alternative opportunities. A basic principle of economic thinking is that opportunity costs and out-of-

pocket costs should be treated alike. Preferences should depend only on relevant differences between options, not on how these differences are labeled. This principle runs counter to the psychological tendencies that make preferences susceptible to superficial variations in form. In particular, a difference that favors outcome A over outcome B can sometimes be framed either as an advantage of A or as a disadvantage of B by suggesting either B or A as the neutral reference point. Because of loss aversion, the difference will loom larger when A is neutral and B-A is evaluated as a loss than when B is neutral and A-B is evaluated as a gain. The significance of such variations of framing has been noted in several contexts.

Thaler (1980) drew attention to the effect of labeling a difference between two prices as a surcharge or a discount. It is easier to forgo a discount than to accept a surcharge because the same price difference is valued as a gain in the former case and as a loss in the latter. Indeed, the credit card lobby is said to insist that any price difference between cash and card purchases should be labeled a cash discount rather than a credit surcharge. A similar idea could be invoked to explain why the price response to slack demand often takes the form of discounts or special concessions (Stigler and Kindahl 1970). Customers may be expected to show less resistance to the eventual cancellation of such temporary arrangements than to outright price increases. Judgments of fairness exhibit the same pattern (Kahneman, Knetsch, and Thaler, in this issue).

Schelling (1981) has described a striking framing effect in a context of tax policy. He points out that the tax table can be constructed by using as a default case either the childless family (as is in fact done) or, say, the modal two-child family. The tax difference between a childless family and a two-child family is naturally framed as an exemption (for the two-child family) in the first frame and as a tax premium (on the childless family) in the second frame. This seemingly innocuous difference has a large effect on judgments of the desired relation between income, family size, and tax. Schelling reported that his students rejected the idea of granting the rich a larger exemption than the poor in the first frame but favored a larger tax premium on the childless rich than on the childless poor in the second frame. Because the exemption and the premium are alternative labels for the same tax differences in the two cases, the judgments violate invariance. Framing the consequences of a public policy in positive or in negative terms can greatly alter its appeal.

The notion of a money illusion is sometimes applied to workers' willingness to accept, in periods of high inflation, increases in nominal wages that do not protect their real income—although they would strenuously resist equivalent wage cuts in the absence of inflation. The essence of the illusion is that, whereas a cut in the nominal wage is

always recognized as a loss, a nominal increase that does not preserve real income may be treated as a gain. Another manifestation of the money illusion was observed in a study of the perceived fairness of economic actions (Kahneman, Knetsch, and Thaler, in press). Respondents in a telephone interview evaluated the fairness of the action described in the following vignette, which was presented in two versions that differed only in the bracketed clauses.

> A company is making a small profit. It is located in a community experiencing a recession with substantial unemployment [but no inflation/and inflation of 12%]. The company decides to [decrease wages and salaries 7%/increase salaries only 5%] this year.

Although the loss of real income is very similar in the two versions, the proportion of respondents who judged the action of the company "unfair" or "very unfair" was 62% for a nominal reduction but only 22% for a nominal increase.

Bazerman (1983) has documented framing effects in experimental studies of bargaining. He compared the performance of experimental subjects when the outcomes of bargaining were formulated as gains or as losses. Subjects who bargained over the allocation of losses more often failed to reach agreement and more often failed to discover a Pareto-optimal solution. Bazerman attributed these observations to the general propensity toward risk seeking in the domain of losses, which may increase the willingness of both participants to risk the negative consequences of a deadlock.

Loss aversion presents an obstacle to bargaining whenever the participants evaluate their own concessions as losses and the concessions obtained from the other party as gains. In negotiating over missiles, for example, the subjective loss of security associated with dismantling a missile may loom larger than the increment of security produced by a similar action on the adversary's part. If the two parties both assign a two-to-one ratio to the values of the concessions they make and of those they obtain, the resulting four-to-one gap may be difficult to bridge. Agreement will be much easier to achieve by negotiators who trade in "bargaining chips" that are valued equally, regardless of whose hand they are in. In this mode of trading, which may be common in routine purchases, loss aversion tends to disappear (Kahneman and Tversky 1984).

IV. The Framing and Weighting of Chance Events

In expected-utility theory, the utility of each possible outcome is weighted by its probability. In prospect theory, the value of an uncertain outcome is multiplied by a decision weight $\pi(p)$, which is a monotonic function of p but is not a probability. The weighting function π

has the following properties. First, impossible events are discarded, that is, $\pi(0) = 0$, and the scale is normalized so that $\pi(1) = 1$, but the function is not well behaved near the end points (Kahneman and Tversky 1979). Second, for low probabilities, $\pi(p) > p$, but $\pi(p) + \pi(1 - p) \leq 1$ (subcertainty). Thus low probabilities are overweighted, moderate and high probabilities are underweighted, and the latter effect is more pronounced than the former. Third, $\pi(pr)/\pi(p) < \pi(pqr)/\pi(pq)$ for all $0 < p, q, r \leq 1$ (subproportionality). That is, for any fixed probability ratio r, the ratio of decision weights is closer to unity when the probabilities are low than when they are high, for example, $\pi(.1)/\pi(.2) > \pi(.4)/\pi(.8)$. A hypothetical weighting function that satisfies these properties is shown in figure 2. Its consequences are discussed in the next section.[3]

Nontransparent Dominance

The major characteristic of the weighting function is the overweighting of probability differences involving certainty and impossibility, for example, $\pi(1.0) - \pi(.9)$ or $\pi(.1) - \pi(0)$, relative to comparable differences in the middle of the scale, for example, $\pi(.3) - \pi(.2)$. In particular, for small p, π is generally subadditive, for example, $\pi(.01) + \pi(.06) > \pi(.07)$. This property can lead to violations of dominance, as illustrated in the following pair of problems.

Problem 7 ($N = 88$). Consider the following two lotteries, described by the percentage of marbles of different colors in each box and the amount of money you win or lose depending on the color of a randomly drawn marble. Which lottery do you prefer?

Option A	90% white	6% red	1% green	1% blue	2% yellow
	\$0	win \$45	win \$30	lose \$15	lose \$15
Option B	90% white	6% red	1% green	1% blue	2% yellow
	\$0	win \$45	win \$45	lose \$10	lose \$15

It is easy to see that option B dominates option A: for every color the outcome of B is at least as desirable as the outcome of A. Indeed, all

3. The extension of the present analysis to prospects with many (nonzero) outcomes involves two additional steps. First, we assume that continuous (or multivalued) distributions are approximated, in the framing phase, by discrete distributions with a relatively small number of outcomes. For example, a uniform distribution on the interval (0, 90) may be represented by the discrete prospect (0, .1; 10, .1; . . . ; 90, .1). Second, in the multiple-outcome case the weighting function, $\pi_p(p_i)$, must depend on the probability vector p, not only on the component p_i, $i = 1, \ldots, n$. For example, Quiggin (1982) uses the function $\pi_p(p_i) = \pi(p_i)/[\pi(p_1) + \ldots + \pi(p_n)]$. As in the two-outcome case, the weighting function is assumed to satisfy subcertainty, $\pi_p(p_1) + \ldots + \pi_p(p_n) \leq 1$, and subproportionality.

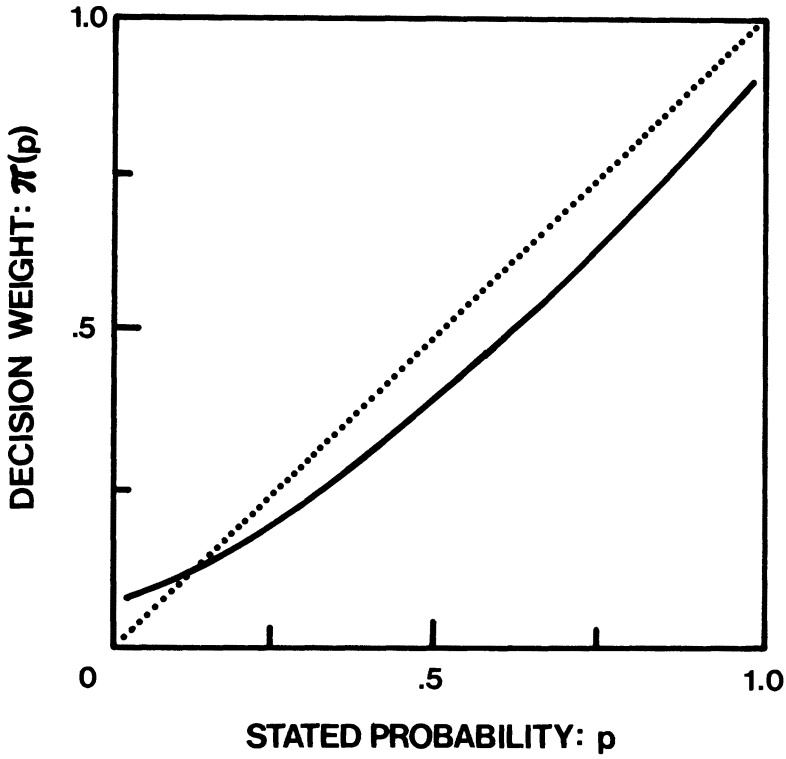

FIG. 2.—A typical weighting function

respondents chose B over A. This observation is hardly surprising because the relation of dominance is highly transparent, so the dominated prospect is rejected without further processing. The next problem is effectively identical to problem 7, except that colors yielding identical outcomes (red and green in B and yellow and blue in A) are combined. We have proposed that this operation is commonly performed by the decision maker if no dominated prospect is detected.

Problem 8 (N = 124). Which lottery do you prefer?

Option C	90% white	6% red	1% green	3% yellow
	\$0	win \$45	win \$30	lose \$15
Option D	90% white	7% red	1% green	2% yellow
	\$0	win \$45	lose \$10	lose \$15

The formulation of problem 8 simplifies the options but masks the relation of dominance. Furthermore, it enhances the attractiveness of

C, which has two positive outcomes and one negative, relative to D, which has two negative outcomes and one positive. As an inducement to consider the options carefully, participants were informed that one-tenth of them, selected at random, would actually play the gambles they chose. Although this announcement aroused much excitement, 58% of the participants chose the dominated alternative C. In answer to another question the majority of respondents also assigned a higher cash equivalent to C than to D. These results support the following propositions. (i) Two formulations of the same problem elicit different preferences, in violation of invariance. (ii) The dominance rule is obeyed when its application is transparent. (iii) Dominance is masked by a frame in which the inferior option yields a more favorable outcome in an identified state of the world (e.g., drawing a green marble). (iv) The discrepant preferences are consistent with the subadditivity of decision weights. The role of transparency may be illuminated by a perceptual example. Figure 3 presents the well-known Müller-Lyer illusion: the top line appears longer than the bottom line, although it is in fact shorter. In figure 4, the same patterns are embedded in a rectangular frame, which makes it apparent that the protruding bottom line is longer than the top one. This judgment has the nature of an inference, in contrast to the perceptual impression that mediates judgment in figure 3. Similarly, the finer partition introduced in problem 7 makes it possible to conclude that option D is superior to C, without assessing their values. Whether the relation of dominance is detected depends on framing as well as on the sophistication and experience of the decision maker. The dominance relation in problems 8 and 1 could be transparent to a sophisticated decision maker, although it was not transparent to most of our respondents.

Certainty and Pseudocertainty

The overweighting of outcomes that are obtained with certainty relative to outcomes that are merely probable gives rise to violations of the expectation rule, as first noted by Allais (1953). The next series of problems (Tversky and Kahneman 1981, p. 455) illustrates the phenomenon discovered by Allais and its relation to the weighting of probabilities and to the framing of chance events. Chance events were realized by drawing a single marble from a bag containing a specified number of favorable and unfavorable marbles. To encourage thoughtful answers, one-tenth of the participants, selected at random, were given an opportunity to play the gambles they chose. The same respondents answered problems 9–11, in that order.

Problem 9 ($N = 77$). Which of the following options do you prefer?
A. a sure gain of \$30 [78%]
B. 80% chance to win \$45 and 20% chance to win nothing [22%]

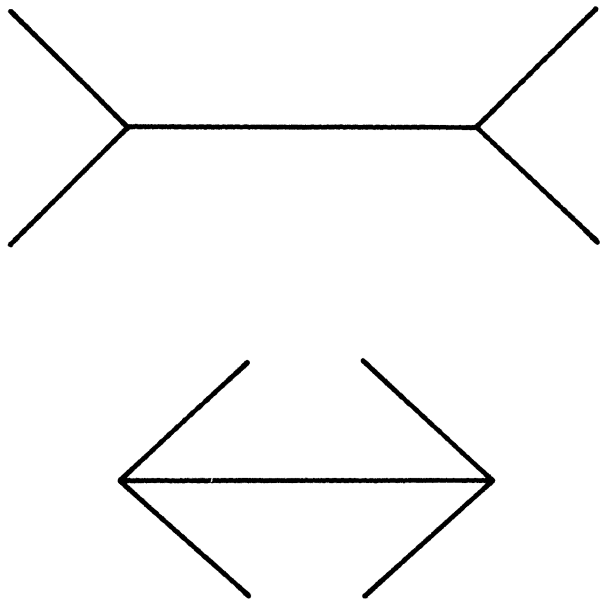

FIG. 3.—The Müller-Lyer illusion

Problem 10 (N = 81). Which of the following options do you prefer?
C. 25% chance to win \$30 and 75% chance to win nothing [42%]
D. 20% chance to win \$45 and 80% chance to win nothing [58%]

Note that problem 10 is obtained from problem 9 by reducing the probabilities of winning by a factor of four. In expected utility theory a preference for A over B in problem 9 implies a preference for C over D in problem 10. Contrary to this prediction, the majority preference switched from the lower prize (\$30) to the higher one (\$45) when the probabilities of winning were substantially reduced. We called this phenomenon the *certainty effect* because the reduction of the probability of winning from certainty to .25 has a greater effect than the corresponding reduction from .8 to .2. In prospect theory, the modal choice in problem 9 implies $v(45)\pi(.80) < v(30)\pi(1.0)$, whereas the modal choice in problem 10 implies $v(45)\pi(.20) > v(30)\pi(.25)$. The observed violation of expected utility theory, then, is implied by the curvature of π (see fig. 2) if

$$\frac{\pi(.20)}{\pi(.25)} > \frac{v(30)}{v(45)} > \frac{\pi(.80)}{\pi(1.0)}.$$

Allais's problem has attracted the attention of numerous theorists, who attempted to provide a normative rationale for the certainty effect by relaxing the cancellation rule (see, e.g., Allais 1979; Fishburn 1982, 1983; Machina 1982; Quiggin 1982; Chew 1983). The following problem

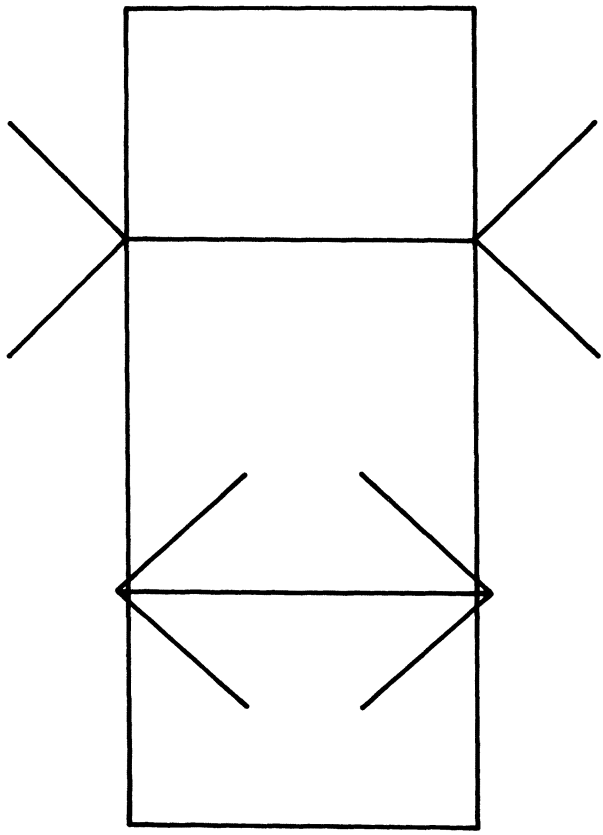

FIG. 4.—A transparent version of the Müller-Lyer illusion

illustrates a related phenomenon, called the *pseudocertainty effect*, that cannot be accommodated by relaxing cancellation because it also involves a violation of invariance.

Problem 11 (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:

E. a sure win of $30 [74%]
F. 80% chance to win $45 and 20% chance to win nothing [26%]

Your choice must be made before the outcome of the first stage is known.

Because there is one chance in four to move into the second stage, prospect E offers a .25 probability of winning $30, and prospect F offers a .25 × .80 = .20 probability of winning $45. Problem 11 is therefore identical to problem 10 in terms of probabilities and out-

comes. However, the preferences in the two problems differ: most subjects made a risk-averse choice in problem 11 but not in problem 10. We call this phenomenon the pseudocertainty effect because an outcome that is actually uncertain is weighted as if it were certain. The framing of problem 11 as a two-stage game encourages respondents to apply cancellation: the event of failing to reach the second stage is discarded prior to evaluation because it yields the same outcomes in both options. In this framing problems 11 and 9 are evaluated alike.

Although problems 10 and 11 are identical in terms of final outcomes and their probabilities, problem 11 has a greater potential for inducing regret. Consider a decision maker who chooses F in problem 11, reaches the second stage, but fails to win the prize. This individual knows that the choice of E would have yielded a gain of $30. In problem 10, on the other hand, an individual who chooses D and fails to win cannot know with certainty what the outcome of the other choice would have been. This difference could suggest an alternative interpretation of the pseudocertainty effect in terms of regret (e.g., Loomes and Sugden 1982). However, the certainty and the pseudocertainty effects were found to be equally strong in a modified version of problems 9–11 in which opportunities for regret were equated across problems. This finding does not imply that considerations of regret play no role in decisions. (For examples, see Kahneman and Tversky [1982, p. 710].) It merely indicates that Allais's example and the pseudocertainty effect are primarily controlled by the nonlinearity of decision weights and the framing of contingencies rather than by the anticipation of regret.[4]

The certainty and pseudocertainty effects are not restricted to monetary outcomes. The following problem illustrates these phenomena in a medical context. The respondents were 72 physicians attending a meeting of the California Medical Association. Essentially the same pattern of responses was obtained from a larger group (N = 180) of college students.

> Problem 12 (N = 72). In the treatment of tumors there is sometimes a choice between two types of therapies: (i) a radical treatment such as extensive surgery, which involves some risk of imminent death,

4. In the modified version—problems 9′–11′—the probabilities of winning were generated by drawing a number from a bag containing 100 sequentially numbered tickets. In problem 10′, the event associated with winning $45 (drawing a number between one and 20) was included in the event associated with winning $30 (drawing a number between one and 25). The sequential setup of problem 11 was replaced by the simultaneous play of two chance devices: the roll of a die (whose outcome determines whether the game is on) and the drawing of a numbered ticket from a bag. The possibility of regret now exists in all three problems, and problem 10′ and 11′ no longer differ in this respect because a decision maker would always know the outcomes of alternative choices. Consequently, regret theory cannot explain either the certainty effect (9′ vs. 10′) or the pseudocertainty effect (10′ vs. 11′) observed in the modified problems.

(ii) a moderate treatment, such as limited surgery or radiation therapy. Each of the following problems describes the possible outcome of two alternative treatments, for three different cases. In considering each case, suppose the patient is a 40-year-old male. Assume that without treatment death is imminent (within a month) and that only one of the treatments can be applied. Please indicate the treatment you would prefer in each case.

Case 1

Treatment A: 20% chance of imminent death and 80% chance of normal life, with an expected longevity of 30 years. [35%]

Treatment B: certainty of a normal life, with an expected longevity of 18 years. [65%]

Case 2

Treatment C: 80% chance of imminent death and 20% chance of normal life, with an expected longevity of 30 years. [68%]

Treatment D: 75% chance of imminent death and 25% chance of normal life, with an expected longevity of 18 years. [32%]

Case 3

Consider a new case where there is a 25% chance that the tumor is treatable and a 75% chance that it is not. If the tumor is not treatable, death is imminent. If the tumor is treatable, the outcomes of the treatment are as follows:

Treatment E: 20% chance of imminent death and 80% chance of normal life, with an expected longevity of 30 years. [32%]

Treatment F: certainty of normal life, with an expected longevity of 18 years. [68%]

The three cases of this problem correspond, respectively, to problems 9–11, and the same pattern of preferences is observed. In case 1, most respondents make a risk-averse choice in favor of certain survival with reduced longevity. In case 2, the moderate treatment no longer ensures survival, and most respondents choose the treatment that offers the higher expected longevity. In particular, 64% of the physicians who chose B in case 1 selected C in case 2. This is another example of Allais's certainty effect.

The comparison of cases 2 and 3 provides another illustration of pseudocertainty. The cases are identical in terms of the relevant outcomes and their probabilities, but the preferences differ. In particular, 56% of the physicians who chose C in case 2 selected F in case 3. The conditional framing induces people to disregard the event of the tumor not being treatable because the two treatments are equally ineffective

in this case. In this frame, treatment F enjoys the advantage of pseudocertainty. It appears to ensure survival, but the assurance is conditional on the treatability of the tumor. In fact, there is only a .25 chance of surviving a month if this option is chosen.

The conjunction of certainty and pseudocertainty effects has significant implications for the relation between normative and descriptive theories of choice. Our results indicate that cancellation is actually obeyed in choices—in those problems that make its application transparent. Specifically, we find that people make the same choices in problems 11 and 9 and in cases 3 and 1 of problem 12. Evidently, people "cancel" an event that yields the same outcomes for all options, in two-stage or nested structures. Note that in these examples cancellation is satisfied in problems that are formally equivalent to those in which it is violated. The empirical validity of cancellation therefore depends on the framing of the problems.

The present concept of framing originated from the analysis of Allais's problems by Savage (1954, pp. 101–4) and Raiffa (1968, pp. 80–86), who reframed these examples in an attempt to make the application of cancellation more compelling. Savage and Raiffa were right: naive respondents indeed obey the cancellation axiom when its application is sufficiently transparent.[5] However, the contrasting preferences in different versions of the same choice (problems 10 and 11 and cases 2 and 3 of problem 12) indicate that people do not follow the same axiom when its application is not transparent. Instead, they apply (nonlinear) decision weights to the probabilities as stated. The status of cancellation is therefore similar to that of dominance: both rules are intuitively compelling as abstract principles of choice, consistently obeyed in transparent problems and frequently violated in nontransparent ones. Attempts to rationalize the preferences in Allais's example by discarding the cancellation axiom face a major difficulty: they do not distinguish transparent formulations in which cancellation is obeyed from nontransparent ones in which it is violated.

V. Discussion

In the preceding sections we challenged the descriptive validity of the major tenets of expected utility theory and outlined an alternative account of risky choice. In this section we discuss alternative theories

5. It is noteworthy that the conditional framing used in problems 11 and 12 (case 3) is much more effective in eliminating the common responses to Allais's paradox than the partition framing introduced by Savage (see, e.g., Slovic and Tversky 1974). This is probably due to the fact that the conditional framing makes it clear that the critical options are identical—after eliminating the state whose outcome does not depend on one's choice (i.e., reaching the second stage in problem 11, an untreatable tumor in problem 12, case 3).

and argue against the reconciliation of normative and descriptive analyses. Some objections of economists to our analysis and conclusions are addressed.

Descriptive and Normative Considerations

Many alternative models of risky choice, designed to explain the observed violations of expected utility theory, have been developed in the last decade. These models divide into the following four classes. (i) Nonlinear functionals (e.g., Allais 1953, 1979; Machina 1982) are obtained by eliminating the cancellation condition altogether. These models do not have axiomatizations leading to a (cardinal) measurement of utility, but they impose various restrictions (i.e., differentiability) on the utility functional. (ii) The expectations quotient model (axiomatized by Chew and MacCrimmon 1979; Weber 1982; Chew 1983; Fishburn 1983) replaces cancellation by a weaker substitution axiom and represents the value of a prospect by the ratio of two linear functionals. (iii) Bilinear models with nonadditive probabilities (e.g., Kahneman and Tversky 1979; Quiggin 1982; Schmeidler 1984; Segal 1984; Yaari 1984; Luce and Narens 1985) assume various restricted versions of cancellation (or substitution) and construct a bilinear representation in which the utilities of outcomes are weighted by a nonadditive probability measure or by some nonlinear transform of the probability scale. (iv) Nontransitive models represent preferences by a bivariate utility function. Fishburn (1982, 1984) axiomatized such models, while Bell (1982) and Loomes and Sugden (1982) interpreted them in terms of expected regret. For further theoretical developments, see Fishburn (1985).

The relation between models and data is summarized in table 1. The stub column lists the four major tenets of expected utility theory. Column 1 lists the major empirical violations of these tenets and cites a few representative references. Column 2 lists the subset of models discussed above that are consistent with the observed violations.

TABLE 1 Summary of Empirical Violations and Explanatory Models

Tenet	Empirical Violation	Explanatory Model
Cancellation	Certainty effect (Allais 1953, 1979; Kahneman and Tversky 1979) (problems 9–10, and 12 [cases 1 and 2])	All models
Transitivity	Lexicographic semiorder (Tversky 1969) Preference reversals (Slovic and Lichtenstein 1983)	Bivariate models
Dominance	Contrasting risk attitudes (problem 2) Subadditive decision weights (problem 8)	Prospect theory
Invariance	Framing effects (Problems 1, 3–4, 5–6, 7–8, 10–11, and 12)	Prospect theory

The conclusions of table 1 may be summarized as follows. First, all the above models (as well as some others) are consistent with the violations of cancellation produced by the certainty effect.[6] Therefore, Allais's "paradox" cannot be used to compare or evaluate competing nonexpectation models. Second, bivariate (nontransitive) models are needed to explain observed intransitivities. Third, only prospect theory can accommodate the observed violations of (stochastic) dominance and invariance. Although some models (e.g., Loomes and Sugden 1982; Luce and Narens 1985) permit some limited failures of invariance, they do not account for the range of framing effects described in this article.

Because framing effects and the associated failures of invariance are ubiquitous, no adequate descriptive theory can ignore these phenomena. On the other hand, because invariance (or extensionality) is normatively indispensable, no adequate prescriptive theory should permit its violation. Consequently, the dream of constructing a theory that is acceptable both descriptively and normatively appears unrealizable (see also Tversky and Kahneman 1983).

Prospect theory differs from the other models mentioned above in being unabashedly descriptive and in making no normative claims. It is designed to explain preferences, whether or not they can be rationalized. Machina (1982, p. 292) claimed that prospect theory is "unacceptable as a descriptive model of behavior toward risk" because it implies violations of stochastic dominance. But since the violations of dominance predicted by the theory have actually been observed (see problems 2 and 8), Machina's objection appears invalid.

Perhaps the major finding of the present article is that the axioms of rational choice are generally satisfied in transparent situations and often violated in nontransparent ones. For example, when the relation of stochastic dominance is transparent (as in the aggregated version of problem 2 and in problem 7), practically everyone selects the dominant prospect. However, when these problems are framed so that the relation of dominance is no longer transparent (as in the segregated version of problem 2 and in problem 8), most respondents violate dominance, as predicted. These results contradict all theories that imply stochastic dominance as well as others (e.g., Machina 1982) that predict the same choices in transparent and nontransparent contexts. The same conclusion applies to cancellation, as shown in the discussion of pseudocertainty. It appears that both cancellation and dominance have normative appeal, although neither one is descriptively valid.

The present results and analysis—particularly the role of transparency and the significance of framing—are consistent with the concep-

6. Because the present article focuses on prospects with known probabilities, we do not discuss the important violations of cancellation due to ambiguity (Ellsberg 1961).

tion of bounded rationality originally presented by Herbert Simon (see, e.g., Simon 1955, 1978; March 1978; Nelson and Winter 1982). Indeed, prospect theory is an attempt to articulate some of the principles of perception and judgment that limit the rationality of choice.

The introduction of psychological considerations (e.g., framing) both enriches and complicates the analysis of choice. Because the framing of decisions depends on the language of presentation, on the context of choice, and on the nature of the display, our treatment of the process is necessarily informal and incomplete. We have identified several common rules of framing, and we have demonstrated their effects on choice, but we have not provided a formal theory of framing. Furthermore, the present analysis does not account for all the observed failures of transitivity and invariance. Although some intransitivities (e.g., Tversky 1969) can be explained by discarding small differences in the framing phase, and others (e.g., Raiffa 1968, p. 75) arise from the combination of transparent and nontransparent comparisons, there are examples of cyclic preferences and context effects (see, e.g., Slovic, Fischhoff, and Lichtenstein 1982; Slovic and Lichtenstein 1983) that require additional explanatory mechanisms (e.g., multiple reference points and variable weights). An adequate account of choice cannot ignore these effects of framing and context, even if they are normatively distasteful and mathematically intractable.

Bolstering Assumptions

The assumption of rationality has a favored position in economics. It is accorded all the methodological privileges of a self-evident truth, a reasonable idealization, a tautology, and a null hypothesis. Each of these interpretations either puts the hypothesis of rational action beyond question or places the burden of proof squarely on any alternative analysis of belief and choice. The advantage of the rational model is compounded because no other theory of judgment and decision can ever match it in scope, power, and simplicity.

Furthermore, the assumption of rationality is protected by a formidable set of defenses in the form of bolstering assumptions that restrict the significance of any observed violation of the model. In particular, it is commonly assumed that substantial violations of the standard model are (i) restricted to insignificant choice problems, (ii) quickly eliminated by learning, or (iii) irrelevant to economics because of the corrective function of market forces. Indeed, incentives sometimes improve the quality of decisions, experienced decision makers often do better than novices, and the forces of arbitrage and competition can nullify some effects of error and illusion. Whether these factors ensure rational choices in any particular situation is an empirical issue, to be settled by observation, not by supposition.

It has frequently been claimed (see, e.g., Smith 1985) that the observed failures of rational models are attributable to the cost of thinking and will thus be eliminated by proper incentives. Experimental findings provide little support for this view. Studies reported in the economic and psychological literature have shown that errors that are prevalent in responses to hypothetical questions persist even in the presence of significant monetary payoffs. In particular, elementary blunders of probabilistic reasoning (Grether 1980; Tversky and Kahneman 1983), major inconsistencies of choice (Grether and Plott 1979; Slovic and Lichtenstein 1983), and violations of stochastic dominance in nontransparent problems (see problem 2 above) are hardly reduced by incentives. The evidence that high stakes do not always improve decisions is not restricted to laboratory studies. Significant errors of judgment and choice can be documented in real world decisions that involve high stakes and serious deliberation. The high rate of failures of small businesses, for example, is not easily reconcilied with the assumptions of rational expectations and risk aversion.

Incentives do not operate by magic: they work by focusing attention and by prolonging deliberation. Consequently, they are more likely to prevent errors that arise from insufficient attention and effort than errors that arise from misperception or faulty intuition. The example of visual illusion is instructive. There is no obvious mechanism by which the mere introduction of incentives (without the added opportunity to make measurements) would reduce the illusion observed in figure 3, and the illusion vanishes—even in the absence of incentives—when the display is altered in figure 4. The corrective power of incentives depends on the nature of the particular error and cannot be taken for granted.

The assumption of the rationality of decision making is often defended by the argument that people will learn to make correct decisions and sometimes by the evolutionary argument that irrational decision makers will be driven out by rational ones. There is no doubt that learning and selection do take place and tend to improve efficiency. As in the case of incentives, however, no magic is involved. Effective learning takes place only under certain conditions: it requires accurate and immediate feedback about the relation between the situational conditions and the appropriate response. The necessary feedback is often lacking for the decisions made by managers, entrepreneurs, and politicians because (i) outcomes are commonly delayed and not easily attributable to a particular action; (ii) variability in the environment degrades the reliability of the feedback, especially where outcomes of low probability are involved; (iii) there is often no information about what the outcome would have been if another decision had been taken; and (iv) most important decisions are unique and therefore provide little opportunity for learning (see Einhorn and Hogarth 1978). The conditions for organizational learning are hardly better. Learning

surely occurs, for both individuals and organizations, but any claim that a particular error will be eliminated by experience must be supported by demonstrating that the conditions for effective learning are satisfied.

Finally, it is sometimes argued that failures of rationality in individual decision making are inconsequential because of the corrective effects of the market (Knez, Smith, and Williams 1985). Economic agents are often protected from their own irrational predilections by the forces of competition and by the action of arbitrageurs, but there are situations in which this mechanism fails. Hausch, Ziemba, and Rubenstein (1981) have documented an instructive example: the market for win bets at the racetrack is efficient, but the market for bets on place and show is not. Bettors commonly underestimate the probability that the favorite will end up in second or third place, and this effect is sufficiently large to sustain a contrarian betting strategy with a positive expected value. This inefficiency is found in spite of the high incentives, of the unquestioned level of dedication and expertise among participants in racetrack markets, and of obvious opportunities for learning and for arbitrage.

Situations in which errors that are common to many individuals are unlikely to be corrected by the market have been analyzed by Haltiwanger and Waldman (1985) and by Russell and Thaler (1985). Furthermore, Akerlof and Yellen (1985) have presented their near-rationality theory, in which some prevalent errors in responding to economic changes (e.g., inertia or money illusion) will (i) have little effect on the individual (thereby eliminating the possibility of learning), (ii) provide no opportunity for arbitrage, and yet (iii) have large economic effects. The claim that the market can be trusted to correct the effect of individual irrationalities cannot be made without supporting evidence, and the burden of specifying a plausible corrective mechanism should rest on those who make this claim.

The main theme of this article has been that the normative and the descriptive analyses of choice should be viewed as separate enterprises. This conclusion suggests a research agenda. To retain the rational model in its customary descriptive role, the relevant bolstering assumptions must be validated. Where these assumptions fail, it is instructive to trace the implications of the descriptive analysis (e.g., the effects of loss aversion, pseudocertainty, or the money illusion) for public policy, strategic decision making, and macroeconomic phenomena (see Arrow 1982; Akerlof and Yellen 1985).

References

Akerlof, G. A., and Yellen, J. 1985. Can small deviations from rationality make significant differences to economic equilibria? *American Economic Review* 75:708–20.

Allais, M. 1953. Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'Ecole Américaine. *Econometrica* 21:503–46.

Allais, M. 1979. The foundations of a positive theory of choice involving risk and a criticism of the postulates and axioms of the American School. In M. Allais and O. Hagen (eds.), *Expected Utility Hypotheses and the Allais Paradox*. Dordrecht: Reidel.

Arrow, K. J. 1982. Risk perception in psychology and economics. *Economic Inquiry* 20:1–9.

Bazerman, M. H. 1983. Negotiator judgment. *American Behavioral Scientist* 27:211–28.

Bell, D. E. 1982. Regret in decision making under uncertainty. *Operations Research* 30:961–81.

Bishop, R. C., and Heberlein, T. A. 1979. Measuring values of extra-market goods: Are indirect measures biased? *American Journal of Agricultural Economics* 61:926–30.

Chew, S. H. 1983. A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. *Econometrica* 51:1065–92.

Chew, S. H., and MacCrimmon, K. 1979. Alpha utility theory, lottery composition, and the Allais paradox. Working Paper no. 686. Vancouver: University of British Columbia.

Clark, H. H., and Clark, E. V. 1977. *Psychology and Language*. New York: Harcourt Brace Jovanovich.

Einhorn, H. J., and Hogarth, R. M. 1978. Confidence in judgment: Persistence of the illusion of validity. *Psychological Review* 85:395–416.

Ellsberg, D. 1961. Risk, ambiguity, and the Savage axioms. *Quarterly Journal of Economics* 75:643–69.

Eraker, S. E., and Sox, H. C. 1981. Assessment of patients' preferences for therapeutic outcomes. *Medical Decision Making* 1:29–39.

Fischhoff, B. 1983. Predicting frames. *Journal of Experimental Psychology: Learning, Memory and Cognition* 9:103–16.

Fishburn, P. C. 1982. Nontransitive measurable utility. *Journal of Mathematical Psychology* 26:31–67.

Fishburn, P. C. 1983. Transitive measurable utility. *Journal of Economic Theory* 31:293–317.

Fishburn, P. C. 1984. SSB utility theory and decision making under uncertainty. *Mathematical Social Sciences* 8:253–85.

Fishburn, P. C. 1985. Uncertainty aversion and separated effects in decision making under uncertainty. Working paper. Murray Hill, N.J.: AT & T Bell Labs.

Fishburn, P. C., and Kochenberger, G. A. 1979. Two-piece von Neumann–Morgenstern utility functions. *Decision Sciences* 10:503–18.

Grether, D. M. 1980. Bayes rule as a descriptive model: The representativeness heuristic. *Quarterly Journal of Economics* 95:537–57.

Grether, D. M., and Plott, C. R. 1979. Economic theory of choice and the preference reversal phenomenon. *American Economic Review* 69:623–38.

Hagen, O. 1979. Towards a positive theory of preferences under risk. In M. Allais and O. Hagen (eds.), *Expected Utility Hypotheses and the Allais Paradox*. Dordrecht: Reidel.

Haltiwanger, J., and Waldman, M. 1985. Rational expectations and the limits of rationality: An analysis of heterogeneity. *American Economic Review* 75:326–40.

Hammond, P. 1985. Consequential behavior in decision trees and expected utility. Institute for Mathematical Studies in the Social Sciences Working Paper no. 112. Stanford, Calif.: Stanford University.

Hansson, B. 1975. The appropriateness of the expected utility model. *Erkenntnis* 9:175–93.

Hausch, D. B.; Ziemba, W. T.; and Rubenstein, M. E. 1981. Efficiency of the market for racetrack betting. *Management Science* 27:1435–52.

Hershey, J. C., and Schoemaker, P. J. H. 1980. Risk taking and problem context in the domain of losses: An expected utility analysis. *Journal of Risk and Insurance* 47:111–32.

Kahneman, D.; Knetsch, J. L.; and Thaler, R. H. In this issue. Fairness and the assumptions of economics.

Kahneman, D.; Knetsch, J. L.; and Thaler, R. In press. Perceptions of fairness: Entitlements in the market. *American Economic Review*.

Kahneman, D., and Tversky, A. 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47:263–91.

Kahneman, D., and Tversky, A. 1982. The psychology of preferences. *Scientific American* 246:160–73.
Kahneman, D., and Tversky, A. 1984. Choices, values, and frames. *American Psychologist* 39:341–50.
Knetsch, J. L., and Sinden, J. A. 1984. Willingness to pay and compensation demanded: Experimental evidence of an unexpected disparity in measures of value. *Quarterly Journal of Economics* 99:507–21.
Knez, P.; Smith, V. L.; and Williams, A. W. 1985. Individual rationality, market rationality and value estimation. *American Economic Review: Papers and Proceedings* 75:397–402.
Loomes, G., and Sugden, R. 1982. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal* 92:805–24.
Luce, R. D., and Krantz, D. H. 1971. Conditional expected utility. *Econometrica* 39:253–71.
Luce, R. D., and Narens, L. 1985. Classification of concatenation measurement structures according to scale type. *Journal of Mathemical Psychology* 29:1–72.
Machina, M. J. 1982. "Expected utility" analysis without the independence axiom. *Econometrica* 50:277–323.
McNeil, B. J.; Pauker, S. G.; Sox, H. C., Jr.; and Tversky, A. 1982. On the elicitation of preferences for alternative therapies. *New England Journal of Medicine* 306:1259–62.
March, J. G. 1978. Bounded rationality, ambiguity, and the engineering of choice. *Bell Journal of Economics* 9:587–608.
Markowitz, H. 1952. The utility of wealth. *Journal of Political Economy* 60:151–58.
Nelson, R. R., and Winter, S. G. 1982. *An Evolutionary Theory of Economic Change*. Cambridge, Mass.: Harvard University Press.
Payne, J. W.; Laughhunn, D. J.; and Crum, R. 1980. Translation of gambles and aspiration level effects in risky choice behavior. *Management Science* 26:1039–60.
Quiggin, J. 1982. A theory of anticipated utility. *Journal of Economic Behavior and Organization* 3:323–43.
Raiffa, H. 1968. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Reading, Mass.: Addison-Wesley.
Russell, T., and Thaler, R. 1985. The relevance of quasi-rationality in competitive markets. *American Economic Review* 75:1071–82.
Savage, L. J. 1954. *The Foundations of Statistics*. New York: Wiley.
Schelling, T. C. 1981. Economic reasoning and the ethics of policy. *Public Interest* 63:37–61.
Schmeidler, D. 1984. Subjective probability and expected utility without additivity. Preprint Series no. 84. Minneapolis: University of Minnesota, Institute for Mathematics and Its Applications.
Schumpeter, J. A. 1954. *History of Economic Analysis*. New York: Oxford University Press.
Segal, U. 1984. Nonlinear decision weights with the independence axiom. Working Paper in Economics no. 353. Los Angeles: University of California, Los Angeles.
Simon, H. A. 1955. A behavioral model of rational choice. *Quarterly Journal of Economics* 69:99–118.
Simon, H. A. 1978. Rationality as process and as product of thought. *American Economic Review: Papers and Proceedings* 68:1–16.
Slovic, P.; Fischhoff, B.; and Lichtenstein, S. 1982. Response mode, framing, and information processing effects in risk assessment. In R. M. Hogarth (ed.), *New Directions for Methodology of Social and Behavioral Science: Question Framing and Response Consistency*. San Francisco: Jossey-Bass.
Slovic, P., and Lichtenstein, S. 1983. Preference reversals: A broader perspective. *American Economic Review* 73:596–605.
Slovic, P., and Tversky, A. 1974. Who accepts Savage's axiom? *Behavioral Science* 19:368–73.
Smith, V. L. 1985. Experimental economics: Reply. *American Economic Review* 75:265–72.
Stigler, G. J., and Kindahl, J. K. 1970. *The Behavior of Industrial Prices*. New York: National Bureau of Economic Research.

Thaler, R. H. 1980. Towards a positive theory of consumer choice. *Journal of Economic Behavior and Organization* 1:39–60.
Tversky, A. 1969. Intransitivity of preferences. *Psychological Review* 76:105–10.
Tversky, A. 1977. On the elicitation of preferences: Descriptive and prescriptive considerations. In D. E. Bell, R. L. Keeney, and H. Raiffa (eds.), *Conflicting Objectives in Decisions*. New York: Wiley.
Tversky, A., and Kahneman, D. 1981. The framing of decisions and the psychology of choice. *Science* 211:453–58.
Tversky, A., and Kahneman, D. 1983. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review* 90:293–315.
von Neumann, J., and Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton, N.J.: Princeton University Press.
Weber, R. J. 1982. The Allais paradox, Dutch auctions, and alpha-utility theory. Working paper. Evanston, Ill.: Northwestern University.
Yaari, M. E. 1984. Risk aversion without decreasing marginal utility. Report Series in Theoretical Economics. London: London School of Economics.