Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice under Risk

CHRIS STARMER[1]

1. *Introduction*

HOW MANY THEORIES of decision making under risk and uncertainty can you think of? Readers of this article will no doubt be familiar with Expected Utility Theory (EUT), the standard theory of individual choice in economics. Many, I expect, will know of a few alternatives to this model. But how many, I wonder, will be aware that these so-called non-expected utility models now number well into double figures? An enormous amount of theoretical effort has been devoted towards developing alternatives to EUT, and this has run hand-in-hand with an ongoing experimental program aimed at testing those theories. The good and proper division of labor suggests that a relatively small group of specialists will be fully aware of the details of this literature. At the same time, the implications of developments in this field are of more than passing interest to the general economist, since what

stimulated developments in non-EU is surely of widespread concern: put bluntly, the standard theory did not fit the facts.

As the standard theory of individual decision making, and as a core component of game theory, EUT constitutes a key building block of a vast range of economic theory. It should be no surprise, therefore, that developing a better understanding of the determinants of individual choice behavior seemed a natural research priority to many theorists. Around two decades of quite intensive research on the topic has generated a great deal of theoretical innovation plus a much richer body of evidence against which models can be judged. There can be few areas in economics that could claim to have sustained such a rich interaction between theory and evidence in an ongoing effort to develop theories in closer conformity with the facts. Considered together, the accumulated theory and evidence present an opportunity to reflect on what has been achieved. Perhaps the most obvious question to address to this literature is this: has it generated, or does it show the prospect of generating, a serious contender for replacing EUT, at least for certain purposes? If the question seems disarmingly straightforward, providing a clear-cut answer will not be.

Identifying a "best theory" naturally requires judgements about the relative importance of predictive accuracy, simplicity, tractability, and so on. Such judgments are complicated by the fact that the evidence, much of which derives from the experimental paradigm, is open to different interpretations.

In what follows, my aim will be to set out what I take to have been key theoretical developments in the area, to review the related evidence and draw conclusions about the current state of play and the prospects for the future. In doing so, rather than simply to present an exhaustive list of models, my aim will be to identify and discuss different modeling *strategies* picking specific models as illustrations. I also intend to narrow my sights in two significant respects. First, my focus will be on *descriptive* as opposed to *normative* issues. Second, I will concentrate on the problem of modeling choices under *risk* as opposed to the more general category of *uncertainty* (the distinction is explained in the next section). Clearing the ground in this way will, I hope, sharpen the focus on one central research problem which continues to motivate much of the research in this arena: the endeavor to develop a "satisfactory" account of *actual* decision behavior in situations of *risk*. It will be a personal view, but one which I hope will help the interested nonspecialist find a trail through this expansive and quite detailed literature.

The paper is organized as follows. Sections 2 and 3 set the scene with discussions of the standard theory and the evidence that prompted theorists to look for alternatives. Section 4 provides the core overview of non-expected utility theories. Section 5 seeks to evaluate what has been achieved so far, and in three subsections I discuss (1) how new theories have fared in a second phase of experimental testing, (2) how new theories may help us to explain a range of phenomena "in the field," and (3) whether non-expected utility theory offers a viable alternative to EUT for everyday theoretical use. In the penultimate Section 6, I discuss two emerging lines of enquiry which I see as particularly exciting paths for future research. A final section offers some concluding reflections.

2. *Where It Began*

Although the primary purpose of this paper is to review alternatives to EUT, that theory provides the natural point of departure, since most of the theories I will be discussing can be understood as generalizations of this base theory.[2] EUT was first proposed by Daniel Bernoulli (1738) in response to an apparent puzzle surrounding what price a reasonable person should be prepared to pay to enter a gamble. It was the conventional wisdom at the time that it would be reasonable to pay anything up to the expected value of a gamble, but Bernoulli presents this counterexample. A coin is flipped repeatedly until a head is produced; if you enter the game, you receive a payoff of, say, $\$2^n$ where n is the number of the throw producing the first head. This is the so-called St. Petersburg game. It is easy to see that its expected monetary payoff is infinite, yet Bernoulli believed most people would only be prepared to pay a relatively small amount to enter it, and he took this intuition as evidence that the "value" of a gamble to an individual is not, in general, equal to its expected monetary value. He proposed a theory in which individuals place subjective values, or "utilities," on monetary outcomes and

[2] I shall not dwell on this account of EUT. For those interested in further discussion an excellent starting place is Paul Schoemaker's (1982) review in this journal.

the value of a gamble is the expectation of these utilities. While Bernoulli's theory—the first statement of EUT—solved the St. Petersburg puzzle, it did not find much favor with modern economists until the 1950s. This is partly explained by the fact that, in the form presented by Bernoulli, the theory presupposes the existence of a cardinal utility scale; an assumption that did not sit well with the drive towards ordinalization during the first half of the twentieth century.

Interest in the theory was revived when John von Neumann and Oskar Morgenstern (1947) showed that the expected utility hypothesis could be derived from a set of apparently appealing axioms on preference. Since then, numerous alternative axiomatizations have been developed, some of which seem highly appealing, some might even say compelling, from a normative point of view (see for example Peter Hammond 1988).[3] To the extent that its axioms can be justified as sound principles of rational choice to which any reasonable person would subscribe, they provide grounds for interpreting EUT *normatively* (as a model of how people ought to choose) and *prescriptively* (as a practical aid to choice). My concern, however, is with how people actually choose, whether or not such choices conform with a priori notions of rationality. Consequently, I will not be delayed by questions about whether particular axioms can or cannot be defended as sound principles of rational choice, and I will start from the presumption that evidence relating to actual behavior should not be discounted purely on the basis that it falls foul of conventional axioms of choice.

For the purpose of understanding al-ternative models of choice, it will be useful to present one set of axioms from which EUT can be derived. In the approach I adopt, at least to begin with, preferences are defined over *prospects* where a prospect is to be understood as a list of consequences with associated probabilities. I will assume throughout that all consequences and probabilities are known to the agent, and hence, in choosing among prospects, the agent can be said to confront a situation of *risk* (in contrast to situations of *uncertainty* in which at least some of the outcomes or probabilities are unknown). I will use lowercase letters in bold (e.g. $\mathbf{q}$, $\mathbf{r}$, $\mathbf{s}$) to represent prospects, and the letter p to represent probabilities (take it that p always lies in the interval $[0,1]$). A given prospect may contain other prospects as consequences, but assuming that such compound prospects can be reduced to simple prospects following the conventional rules of probability, any prospect $\mathbf{q}$ can be represented by a probability distribution $\mathbf{q} = (p_1, \ldots, p_n)$ over a fixed set of pure consequences $X = (x_1, \ldots, x_n)$ where p_i is the probability of x_i, $p_i \geq 0$ for all i, and $\Sigma_i p_i = 1$. Hence, the elements of X are to be understood as an exhaustive and mutually exclusive list of possible consequences which may follow from a particular course of action. While this notation allows a prospect to be written simply as vector of probabilities (as $\mathbf{q}$ above) it will sometimes be useful to be explicit about the consequences too (e.g. by writing $\mathbf{q} = (x_1, p_1; \ldots; x_n, p_n)$).

Given these preliminaries, the expected utility hypothesis can be derived from three axioms: ordering, continuity, and independence. The ordering axiom requires both completeness and transitivity. Completeness entails that for all $\mathbf{q}$, $\mathbf{r}$: *either* $\mathbf{q} \geq \mathbf{r}$ *or* $\mathbf{r} \geq \mathbf{q}$ *or both* where $\geq$ represents the relation "is (weakly) preferred to." Transitivity requires that

---

[3] Such arguments, whilst widely accepted, are nevertheless controversial. See, for example, Paul Anand (1992) and Sugden (1991).

for all **q**, **r**, **s**: *if* **q** ≽ **r** *and* **r** ≽ **s**, *then* **q** ≽ **s**. Continuity requires that for all prospects **q**, **r**, **s** where **q** ≽ **r** and **r** ≽ **s**: *there exists some p such that* (**q**, *p*; **s**, $1 - p$) ~ **r**, where ~ represents the relation of indifference and (**q**, *p*; **s**, $1 - p$) represents a (compound) prospect which results in **q** with probability *p*; **s** with probability $1 - p$. Together the axioms of ordering and continuity imply that preferences over prospects can be represented by a function *V(.)* which assigns a real-valued index to each prospect. The function *V(.)* is a representation of preference in the sense that $V(\mathbf{q}) \geq V(\mathbf{r}) \Leftrightarrow \mathbf{q} \succcurlyeq \mathbf{r}$: that is, an individual will choose the prospect **q** over the prospect **r** if, and only if, the value assigned to **q** by *V(.)* is no less than that assigned to **r**.

To assume the existence of some such preference function has seemed, to many economists, the natural starting point for any economic theory of choice; it amounts to assuming that agents have well-defined preferences, while imposing minimal restriction on the precise form of those preferences. For those who endorse such an approach, the natural questions center around what further restrictions can be placed on *V(.)*. The independence axiom of EUT places quite strong restrictions on the precise form of preferences: it is this axiom which gives the standard theory most of its empirical content (and it is the axiom which most alternatives to EUT will relax). Independence requires that for all prospects **q**, **r**, **s**: *if* **q** ≽ **r** *then* (**q**, *p*; **s**, $1 - p$) ≽ (**r**, *p*; **s**, $1 - p$), for all *p*. If all three axioms hold, preferences can be represented by:

$$V(\mathbf{q}) = \Sigma_i p_i \cdot u(x_i) \qquad (1)$$

where **q** is any prospect, and *u(.)* is a "utility" function defined on the set of consequences.

The concept of risk is pervasive in economics, so economists naturally need a theory of individual decision making under risk. EUT has much to recommend itself in this capacity. The theory has a degree of intuitive appeal. It seems almost trivially obvious that any satisfactory theory of decision making under risk will necessarily take account of both the consequences of choices and their associated probabilities. These are, by definition, the dimensions relevant in the domain of risk. EUT provides one very simple way of combining probabilities and consequences into a single "measure of value" which has a number of appealing properties. One such property is monotonicity, which can be defined as follows. Let $x_1, \ldots, x_n$ be consequences ordered from worst (x_1) to best (x_n). We may say that one prospect $\mathbf{q} = (p_{q1}, \ldots, p_{qn})$ first-order stochastically dominates another prospect $\mathbf{r} = (p_{r1}, \ldots, p_{rn})$ if for all $i = 1, \ldots, n$:

$$\sum_{j=i}^{n} p_{qj} \geq \sum_{j=i}^{n} p_{rj} \qquad (2)$$

with a strict inequality for at least one *i*. Monotonicity is the property that stochastically dominating prospects are preferred to prospects which they dominate, and it is widely held that any satisfactory theory—descriptive or normative—should embody monotonicity. I will have more to say about this later.

The shape of the utility function also has a simple behavioral interpretation whereby concavity (convexity) of *u(.)* implies risk averse (prone) behavior; an agent with a concave utility function will always prefer a certain amount *x* to any risky prospect with expected value equal to *x*. Modeling risk preferences in this way does collapse some potentially distinct concepts into a single function: any attitude to chance (e.g., like or dislike of taking risks) and any attitude

towards consequences (e.g. a diminishing marginal utility of money) must all be captured by the utility function. That need not imply any weakness of the theory. Indeed it is precisely the simplicity and economy of EUT that has made it such a powerful and tractable modeling tool. My concern, however, is with the descriptive merits of the theory and, from this point of view, a crucial question is whether EUT provides a sufficiently accurate representation of actual choice behavior. The evidence from a large number of empirical tests has raised some real doubts on this score.

3. *Descriptive Limitations of Expected Utility Theory—The Early Evidence*

Empirical studies dating from the early 1950s have revealed a variety of patterns in choice behavior that appear inconsistent with EUT. I shall not attempt a full-blown review of this evidence.[4] Instead, I discuss one or two examples to illustrate the general nature of this evidence, and offer a discussion of its role in stimulating the development of new theories. With hindsight, it seems that violations of EUT fall under two broad headings: those which have possible explanations in terms of some "conventional" theory of preferences and those which apparently do not. The former category consists primarily of a series of observed violations of the independence axiom of EUT; the latter of evidence that seems to challenge the assumption that choices derive from well-defined preferences. Let us begin with the former.

There is now a large body of evidence indicating that actual choice behavior may *systematically* violate the independence axiom. Two examples of such phenomena, first discovered by Maurice Allais (1953), have played a particularly important role in stimulating and shaping theoretical developments in non-EU theory. These are the so-called *common consequence effects* and *common ratio effects*. The first sighting of such effects came in the form of the following pair of hypothetical choice problems. In the first you have to imagine choosing between the two prospects: $s_1 = (\$1M, 1)$ or $r_1 = (\$5M, 0.1; \$1M, 0.89; 0, 0.01)$. The first option gives one million U.S. dollars for sure; the second gives five million with a probability of 0.1; one million with a probability of 0.89, otherwise nothing.[5] What would you choose? Now consider a second problem where you have to choose between the two prospects: $s_2 = (\$1M, 0.11; 0, 0.89)$ or $r_2 = (\$5M, 0.1; 0, 0.9)$. What would you do if you really faced this choice?

Allais believed that EUT was not an adequate characterization of individual risk preferences and he designed these problems as a counterexample. As we shall shortly see, a person with expected utility preferences would either choose both "s" options, or choose both "r" options across this pair of problems. Allais expected that people faced with these choices might opt for s_1 in the first problem, lured by the certainty of becoming a millionaire, and select r_2 in the second choice where the odds of winning seem very similar, but the prizes very different. Evidence quickly emerged that many people did respond to these problems as Allais had predicted. This is the famous "Allais paradox" and it is one example of the more general common consequence effect.

Most examples of the common consequence effect have involved choices between pairs of prospects of the following

[4] Those interested in more thorough reviews are recommended to consult Schoemaker (1982) and, more recently, Colin Camerer (1995).

[5] In Allais' original examples, consequences were French Francs.

form: $\mathbf{s}^* = (y, p; c, 1-p)$ and $\mathbf{r}^* = (\mathbf{q}, p;$ $c, 1-p)$, where $\mathbf{q} = (x, \lambda; 0, 1-\lambda)$ and $0 < \lambda < 1$.[6] The payoffs c, x and y are nonnegative (usually monetary) consequences such that $x > y$. Notice that both prospects $\mathbf{s}^*$ and $\mathbf{r}^*$ give outcome c with probability $1-p$: this is the "common consequence" and it is an obvious implication of the independence axiom of EUT that choices between $\mathbf{s}^*$ and $\mathbf{r}^*$ should be independent of the value of c.[7] Numerous studies, however, have found that choices between prospects with this basic structure are systematically influenced by the value of c. More specifically, a variety of experimental studies[8] reveal a tendency for individuals to choose $\mathbf{s}^*$ when $c = y$, and $\mathbf{r}^*$ when $c = 0$.

A closely related phenomenon, also discovered by Allais, is the so-called *common ratio effect*. Suppose you had to make a choice between $3000 for sure, or entering a gamble with an 80 percent chance of getting $4000 (otherwise nothing). What would you choose? Now think about what you would do if you had to choose either a 25 percent chance of gaining $3000 or a 20 percent chance of gaining $4000. A good deal of evidence suggests that many people would opt for the certainty of $3000 in the first choice and opt for the 20 percent chance of $4000 in the second. Such a pattern of choice, however, is inconsistent with EUT and would constitute one example of the common ratio effect. More generally, this phenomenon is observed in choices among pairs of problems with the following form:

$\mathbf{s}^{**} = (y, p; 0, 1-p)$ and $\mathbf{r}^{**} = (x, \lambda p; 0, 1-\lambda p)$ where $x > y$. Assume that the ratio of "winning" probabilities (λ) is constant, then for pairs of prospects of this structure, EUT implies that preferences should not depend on the value of p,[9] yet numerous studies[10] reveal a tendency for individuals to switch their choice from $\mathbf{s}^{**}$ to $\mathbf{r}^{**}$ as p falls.

It would, of course, be unrealistic to expect any theory of human behavior to predict accurately one hundred percent of the time. Perhaps the most one could reasonably expect is that departures from such a theory be equally probable in each direction. These phenomena, however, involve systematic (i.e., predictable) directions in majority choice. As evidence against the independence axiom accumulated, it seemed natural to wonder whether assorted violations of it might be revealing some underlying feature of preferences that, if properly understood, could form the basis of a unified explanation. Consequently, a wave of theories designed to explain the evidence began to emerge at the end of the 1970s. Most of these theories have the following features in common: (i) preferences are represented by some function $V(.)$ defined over individual prospects; (ii) the function satisfies ordering and continuity; and (iii) while $V(.)$ is designed to permit observed violations of the independence axiom, the principle of monotonicity is retained. I will call theories with these properties

---

[6] It will be convenient to use a scaling factor λ at several points in the paper, so to avoid repetition, assume $0 < \lambda < 1$ throughout.

[7] The original Allais problems are recovered from this generalization setting $x = \$5M$; $y = \$1M$, $p = 0.11$ and $\lambda = 10/11$.

[8] Examples include Herbert Moskowitz (1974), Paul Slovic and Amos Tversky (1974), Kenneth MacCrimmon and Stig Larsson (1979).

[9] To see why, consider any pair of options ($\mathbf{s}_1^{**}$, $\mathbf{r}_1^{**}$) where $p = p_1$, then define a further pair of options ($\mathbf{s}_2^{**}$, $\mathbf{r}_2^{**}$) identical except having a *lower* value of $p = p_2$. Since there must be some α, ($1 > \alpha > 0$), such that $p_2 = \alpha p_1$, we can write $\mathbf{s}_2^{**} = (\mathbf{s}_1^{**}, \alpha; 0, 1-\alpha)$ and $\mathbf{r}_2^{**} = (\mathbf{r}_1^{**}, \alpha; 0, 1-\alpha)$. It then follows directly from independence that choices between such pairs of prospects should not depend on the value of p.

[10] Examples include Loomes and Sugden (1987), Starmer and Sugden (1989), and Raymond Battalio, John Kagel, and Komain Jiranyakul (1990).

conventional theories. The general spirit of the approach is to seek "well behaved" theories of preference consistent with observed violations of independence; I call this general approach the *conventional strategy.*

There is evidence to suggest that failures of EUT may run deeper than violations of independence. Two assumptions implicit in any conventional theory are *procedure invariance* (preferences over prospects are independent of the method used to elicit them) and *description invariance* (preferences over prospects are purely a function of the probability distributions of consequences implied by prospects and *do not depend on how those given distributions are described*). While these assumptions probably seem natural to most economists—so natural that they are rarely even discussed when stating formal theories—there is ample evidence that, in practice, both assumptions fail.

One well-known phenomenon, often interpreted as a failure of procedure invariance, is *preference reversal*. The classic preference reversal experiment requires individuals to carry out two distinct tasks (usually separated by some other intervening tasks). The first task requires the subject to choose between two prospects: one prospect (often called the $-bet) offers a small chance of winning a "good" prize; the other (the "P-bet") offers a larger chance of winning a smaller prize. The second task requires the subject to assign monetary values—usually minimum selling prices denoted $M(\$)$ and $M(P)$—to the two prospects. Repeated studies[11] have revealed a tendency for individuals to chose the P-bet (i.e., reveal $P > \$$) while placing a higher value on the $-bet (i.e., $M(\$) > M(P)$). This is

the so-called *preference reversal phenomenon* first observed by psychologists Sarah Lichtenstein and Paul Slovic (1971) and Harold Lindman (1971). It presents a puzzle for economics because, viewed from the standard theoretical perspective, both tasks constitute ways of asking essentially the same question, that is, "which of these two prospects do you prefer?" In these experiments, however, the ordering revealed appears to depend upon the elicitation procedure.

One explanation for preference reversal suggests that choice and valuation tasks may invoke different mental processes which in turn generate different orderings of a given pair of prospects (see Slovic 1995). Consequently, the rankings observed in choice and valuation tasks cannot be explained with reference to a *single* preference ordering. An alternative interpretation explains preference reversal as a failure of transitivity (see Graham Loomes and Robert Sugden 1983): assuming that the valuation task reveals true monetary valuations, (i.e., $M(\$) \sim \$$; $M(P) \sim P$), preference reversal implies $P > \$ \sim M(\$) > M(P) \sim P$; which involves a violation of transitivity (assuming that more money is preferred to less). Although attempts have been made to explain the evidence in ways which preserve conventional assumptions—see for example Charles Holt (1986); Edi Karni and Zvi Safra (1987); Uzi Segal (1988)—the weight of evidence suggests that failures of transitivity and procedure invariance both contribute to the phenomenon (Loomes, Starmer, and Sugden 1989; Tversky, Slovic, and Daniel Kahneman 1990).

There is also widespread evidence that very minor changes in the presentation or "framing" of prospects can have dramatic impacts upon the choices of decision makers: such effects are

---

[11] Reviews of this evidence are contained in Tversky and Richard Thaler (1990), Daniel Hausman (1992), and Timo Tammi (1997).

failures of description invariance. Here is one famous example due to Tversky and Kahneman (1981) in which two groups of subjects—call them groups I and II—were presented with the following cover story:

> *"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 estimate of the consequences of the programs are as follows:"*

Each group then faced a choice between two policy options.

> Options presented to group I:
> *"If program A is adopted, 200 people will be saved.*
>
> *If program B is adopted, there is a 1/3 probability that 600 people will be saved, and a 2/3 probability that no people will be saved."*
>
> Options presented to group II:
> *"If program C is adopted, 400 people will die.*
>
> *If program D is adopted, there is a 1/3 probability that nobody will die, and a 2/3 probability that 600 people will die."*

The two pairs of options are stochastically equivalent. The only difference is that the group I description presents the information in terms of *lives saved* while the information presented to group II is in terms of lives lost. Tversky and Kahneman found a very striking difference in responses to these two presentations: 72 percent of subjects preferred option A to option B while only 22 percent of subjects preferred C to D. Similar patterns of response were found amongst groups of undergraduate students, university faculty, and practicing physicians.

Failures of procedure invariance and description invariance appear, on the face of it, to challenge the very idea

that choices can, in general, be represented by *any* well behaved preference function. If that is right, they lie outside the explanatory scope of the conventional strategy. Some might even be tempted to say they lie outside the scope of economic theory altogether. That stronger claim, however, is controversial, and I will not be content to put away such challenging evidence so swiftly. For present purposes, suffice it to make two observations. First, whether or not we have adequate economic theories of such phenomenon, the "Asian disease" example is clearly suggestive that framing effects have a bearing on issues of genuine economic relevance. Second, there are at least some theories of choice that predict phenomena like preference reversal and framing effects, and some of these models have been widely discussed in the economics literature. Although most of these theories—or at least the ones I will discuss—draw on ideas about preference to explain choices, they do so in unorthodox ways, and many draw on concepts more familiar to psychologists than economists. The one feature common to this otherwise heterodox bunch of theories is that none of them can be reduced to, or expressed purely in terms of, a single preference function $V(.)$ defined over individual prospects. I will call such models *nonconventional theories*. These theories step into what has been relatively uncharted water for the economics profession. One of the aims of this piece is to reflect on the relative merits of the conventional and nonconventional approaches.

4. *Non-Expected Utility Theories*

4.1 *The Conventional Strategy*

One way to approach this literature is by asking a question that motivated a number of theories: what properties would a conventional theory of preference need
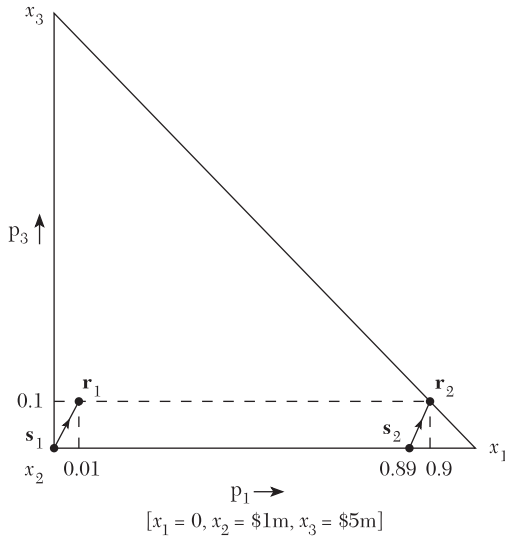
Figure 1a. The Allais Paradox Problems in a Probability Triangle

Figure 1b. Expected Utility Indifference Curves

to have in order to explain the known violations of independence? In order to pursue that question, it will be helpful to introduce an expositional device known as the probability triangle diagram,[12] and this will also prove useful as a vehicle for comparing the predictions of alternative theories.

Consider the class of prospects defined over three outcomes x_1, x_2, x_3 such that $x_1 < x_2 < x_3$. Since any such prospects can be described as a vector of probabilities $(p_1,\ 1 - p_1 - p_3,\ p_3)$ we can also locate them, graphically, in two-dimensional probability space. Figure 1a is a probability triangle that does this for the four prospects $\{s_1, r_1, s_2, r_2\}$ from the original Allais paradox problems. By convention, the horizontal axis measures the probability of the worst consequence ($\$0$) increasing from left to right; the vertical axis measures the probability of the best consequence

[12] Although the probability triangle had appeared in the literature many years before (see Jacob Marschak 1950) Mark Machina's use of it in the 1980s (see below) popularized it to the extent that some have called this diagram the "Machina triangle."
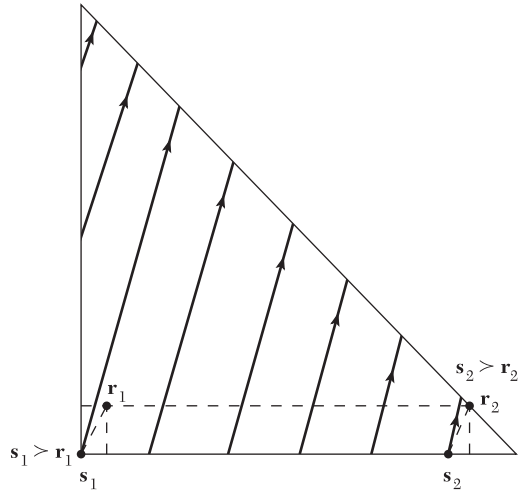
($\$5M$) increasing from bottom to top. Hence s_1, which results in the intermediate consequence of $\$1M$ for sure, is located at the bottom left corner of the triangle; s_2 and r_2, which each assign positive probability to only two of the three possible consequences, are located on the triangle boundaries; while r_1, which assigns positive probability to all three consequences, lies on the interior of the triangle. Two lines have been drawn in the triangle joining the pairs of prospects involved in the two choices. It is easy to establish that these two lines are parallel.

Given ordering plus continuity, preferences over prospects in any given triangle can be represented by a set of indifference curves, hence, every conventional theory implies the existence of a set of indifference curves in this space though the precise form of indifference curves varies between them.

The addition of the independence axiom of EUT restricts the set of indifference curves to being *upward sloping* (left to right) *linear* and *parallel*. One such set of indifference curves is illustrated in Figure 1b (preferences are increasing moving northwest). Independence is a
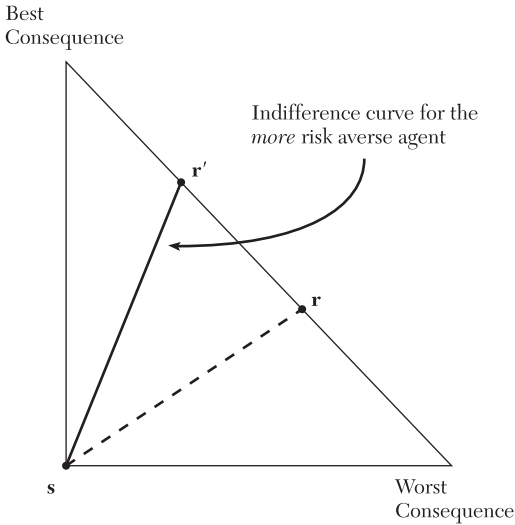
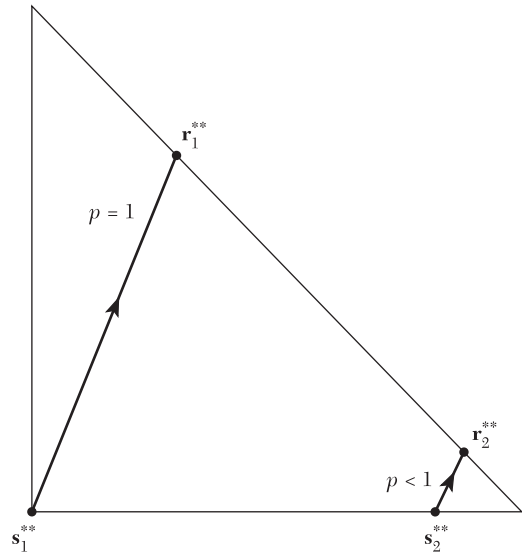*Figure* 1c. Different Degrees of Risk Aversion in EUT

Figure 2. Common Ratio Prospects

strong restriction which leaves only one feature of the indifference curves undetermined, that is, their slope. In EUT, the slope of the indifference curves reflects attitude to risk and may vary between individuals: the more risk averse the individual, the steeper the slope of their indifference curves. To see why, look at Figure 1c and consider two individuals: person 1 has indifference curves with the slope of the dashed line (hence $\mathbf{s} \sim \mathbf{r}$); person 2 has indifference curves with the slope of the solid line (hence $\mathbf{s} \sim \mathbf{r}'$). Person 2 can be seen to be the more risk averse in the sense that, as we move northwest along the hypotenuse, relative to person 1, we must give her a higher chance of winning the best outcome in the riskier prospect in order to generate indifference with the safe prospect $\mathbf{s}$.

In relation to the Allais paradox problems in Figure 1b, for a given individual, EUT allows three possibilities. Indifference curves could have a steeper slope than the lines connecting prospects, in which case $\mathbf{s}_1 > \mathbf{r}_1$ and $\mathbf{s}_2 > \mathbf{r}_2$. This is the case represented in Figure 1b. Alternatively, indifference curves

could have a less steep slope (in which case $\mathbf{r}_1 > \mathbf{s}_1$ and $\mathbf{r}_2 > \mathbf{s}_2$). Finally, the slope of indifference curves could correspond exactly with that of the lines joining pairs of prospects, in which case $\mathbf{r}_1 \sim \mathbf{s}_1$ and $\mathbf{r}_2 \sim \mathbf{s}_2$. But as noted above, people often violate EUT, revealing $\mathbf{s}_1 > \mathbf{r}_1$ in the left-hand problem, $\mathbf{r}_2 > \mathbf{s}_2$ in the right-hand problem. Relative to the predictions of EUT, in choosing $\mathbf{r}_2$ over $\mathbf{s}_2$ these people are being more risk seeking than they should be, given their choice of $\mathbf{s}_1$ over $\mathbf{r}_1$.

A similar tendency is apparent in the common ratio effect. A pair of common ratio problems is illustrated in Figure 2. The pair of prospects $\{\mathbf{s}_1^{**}, \mathbf{r}_1^{**}\}$, near the left edge of the triangle, correspond with the common ratio problems where $p = 1$. As p falls, we generate pairs of prospects like $\{\mathbf{s}_2^{**}, \mathbf{r}_2^{**}\}$, located on parallel lines further to the right in the triangle. Assuming expected utility preferences, an individual must either prefer the "safer option" in *both* choices *or* the "riskier option" in *both* choices, yet many people choose $\mathbf{s}_1^{**}$ over $\mathbf{r}_1^{**}$ and $\mathbf{r}_2^{**}$ over $\mathbf{s}_2^{**}$. This is the common ratio

effect and, as in the common consequence effect, relative to the predictions of EUT, there is an "inconsistency" in the risk attitudes revealed across their choices.

Viewed in the context of the triangle, this inconsistency is suggestive of a systematic pattern: relative to the predictions of EUT, choices between prospects located in the bottom right-hand corner appear more risk prone than should be expected given preferences revealed for choices located leftwards and/or upwards in the triangle. Any conventional theory seeking to explain these standard violations of EUT will therefore need at least one quite specific property: indifference curves determining preferences over pairs of prospects located near the right-hand corner of a given triangle—like, say $\{s_2^{**}, r_2^{**}\}$—will need to be relatively flat (reflecting more risk-prone behavior), compared with indifference curves determining choices over pairs of prospects, like $\{s_1^{**}, r_1^{**}\}$, near to the left-hand edge of the triangle. All of the proposed conventional alternatives to EUT are able to generate this property, though they do so in a variety of ways.

4.1.1 *The "Fanning-out" Hypothesis*

Having observed this apparent connection between different violations of independence, Mark Machina (1982) proposed an analytical extension of EUT (termed "generalized expected utility analysis"), along with a specific hypothesis on the shape of non-expected utility indifference curves. Analytically, he noted that under expected utility, where $V(\mathbf{q}) = \Sigma\ U(x_i) \cdot p_i$, the utility values $U(x_i) = \delta V(\mathbf{q})/\delta p_i$ are the probability derivatives of $V(.)$. He then showed that standard expected utility results (e.g., risk aversion $\Leftrightarrow$ concavity of $U(.)$) also hold for the probability derivatives $U(x_i;\mathbf{q}) = \delta V(\mathbf{q})/\delta p_i$ of smooth non-
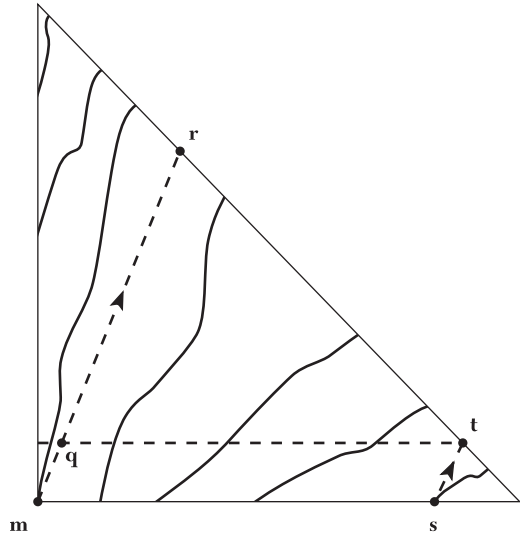

Figure 3. Indifference Curves in Generalized Expected Utility

expected utility preference functions $V(.)$, so that $U(.;\mathbf{q})$ can be thought of as the "local utility function" of $V(.)$ about $\mathbf{q}$. For example, the property "concavity of $U(.;\mathbf{q})$ at every $\mathbf{q}$" is equivalent to global risk aversion of $V(.)$.

Given the existence of phenomena like the common ratio and common consequence effects, Machina hypothesized that the local utility functions $U(.;\mathbf{q})$ become more concave as we move from (first order) stochastically dominated to stochastically dominating distributions. Loosely speaking, this essentially empirical assumption (which Machina calls "Hypothesis II") implies a tendency for agents to become more risk averse as the prospects they face get better; in the context of the triangle, it means that indifference curves become steeper, or "fan out" as we move northwest. Figure 3 illustrates the general pattern of indifference curves implied by Hypothesis II. Notice that they are drawn as wavy lines: generalized expected utility theory requires indifference curves to be smooth but does not imply that they must be linear

(though they may be). It is very easy to see that this fanning-out property generates implications consistent with the common consequence and common ratio effects. Since indifference curves are relatively steeply sloped in the neighborhood of prospect **m**, **m** lies on a higher indifference curve than **q** or **r**. Flatter indifference curves in the bottom right-hand corner of the triangle are such that **t** lies on a higher indifference curve than **s**. Hence, for an individual whose indifference curves fan out we can construct prospects over which we will observe a common consequence effect (e.g. **m** > **q** and **t** > **s**) and a common ratio effect (e.g. **m** > **r** and **t** > **s**).

A whole family of models have this fanning-out property and, within this family, one important subset consists of those models that restrict indifference curves to be linear. One example is Soo Hong Chew and MacCrimmon's (1979) weighted utility theory in which preferences over prospects are represented by the function:

$$V(\mathbf{q}) = [\Sigma p_i \cdot g(x_i) \cdot u(x_i)] / [\Sigma p_i \cdot g(x_i)] \quad (3),$$

where $u(.)$ and $g(.)$ are two different functions assigning non-zero weights to all consequences. The model incorporates EUT as the special case in which the weights assigned by $g(.)$ are identical for every consequence. Weighted utility has been axiomatized by, among others, Chew and MacCrimmon (1979), Chew (1983), and Peter Fishburn (1983), and different variants are discussed in Fishburn (1988). Essentially these axiomatizations involve a weakened form of the independence axiom which constrains indifference curves to be linear without requiring them to be parallel. One version of weak independence is this: *if* **q** > **r** *then for each* p_q *there exists a corresponding* p_r *such that* (**q**, p_q; **s**, $1 - p_q$) > (**r**, p_r; **s**, $1 - p_r$) *for all* **s**. If we think in terms of preferences in the triangle dia-
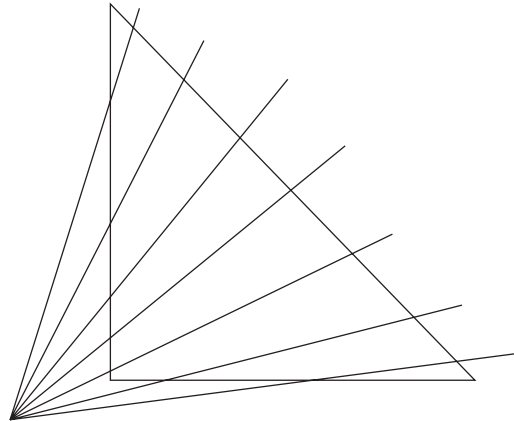

Figure 4. Weighted Utility Theory with Indifference Curves Fanning Out

gram, excepting the special case of EUT, this axiom has the effect of requiring there to be some point at which all indifference curves cross. The location of this point, which could lie inside or outside of the triangle boundary, depends upon the specifications of the functions $u(.)$, and $g(.)$. Transitivity can be preserved by making the point from which curves radiate lie outside the boundary of the triangle and, in order to explain the common ratio and common consequence effects, the origin of indifference curves must lie somewhere to the southwest of the triangle, as in Figure 4. Having restricted the model in this way,[13] we can then understand it as a special case of Machina's theory (including Hypothesis II) in which indifference curves are constrained to be linear.

It is not obvious to me that weak independence has much, if any, intuitive appeal, and the main rationale for assuming it in weighted utility theory is presumably that it results in a simple mathematical function capable of generating fanning out and hence explaining the early violations of EUT. Other models with very similar properties

[13] Chew and MacCrimmon (1979b) explain the conditions necessary to generate this property.

have been based on psychologically grounded hypotheses. One example is the theory of disappointment developed by David Bell (1985) and Loomes and Sugden (1986). While this theory lacks axiomatic foundations, it has a more obvious intuitive interpretation. In the version presented by Loomes and Sugden, preferences over prospects can be represented by the function:

$$V(\mathbf{q}) = \Sigma_i p_i[u(x_i) + D(u(x_i) - \underline{U})] \quad (4),$$

where $u(x_i)$ is interpreted as a measure of "basic" utility (that is, the utility of x_i, considered in isolation from the other consequences of $\mathbf{q}$) and $\underline{U}$ is a measure of the "prior expectation" of the utility from the prospect. The model assumes that if the outcome of a prospect is worse than expected (i.e., if $u(x_i) < \underline{U}$) a sense of disappointment will be generated. On the other hand, an outcome better than expected will stimulate "elation." With $D(.) = 0$, the model reduces to EUT. This additional function, however, is intended to capture a particular intuition about human psychology: that people dislike disappointment and so act to avoid it. More specifically, this is captured by assuming that agents are "disappointment averse" ($D(h)$ is concave for $h < 0$) and "elation prone" ($D(.)$ is convex for $h > 0$). The theory then implies a tendency for indifference curves to fan out in the triangle. The theory of disappointment has close affinity with earlier models based on *moments of utility*. In EUT, the value of a prospect is the (probability weighted) mean of utility. Allais (1979) proposed a model in which $V(.)$ may also depend on the second moment of utility, that is, the variance of utility about the mean. Hagen (1979) extended this idea to include the third moment of utility, or skewness. Sugden(1986) shows that properties of $D(.)$ imposed in disappointment theory can be interpreted as restrictions on Hagen's general model of moments.

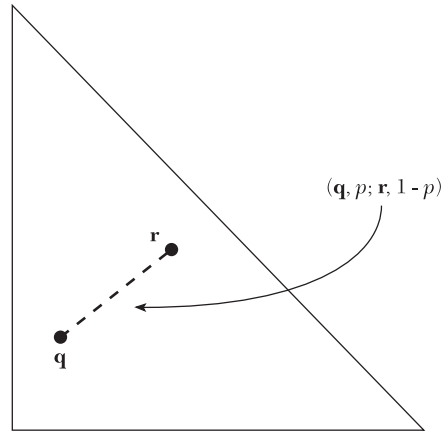*Figure* 5a. Probability Mixtures of Prospects **q**, **r**

A series of other models with linear indifference curves including implicit expected utility (Eddie Dekel 1986) and implicit weighted utility (Chew 1989) allow fanning out, but also permit more complex patterns. For example, Faruk Gul (1991) and William Neilson (1992) present models based on implicit expected utility which generate a mixture of fanning-in and fanning-out within a given triangle.[14] The crucial axiom in these models is a weakened form of independence called *betweenness*: if $\mathbf{q} > \mathbf{r}$, then $\mathbf{q} > (\mathbf{q}, p; \mathbf{r}, (1-p)) > \mathbf{r}$ for all $p < 1$. It is this assumption that imposes linearity on indifference curves and, conversely, it is implied by any model that assumes linear indifference curves.

Behaviorally, betweenness implies that any probability mixture of two lotteries will be ranked between them in terms of preference and, given continuity, an individual will be indifferent to randomization among equally valued prospects. To understand the connection between these behavioral and geometric properties, look at Figure 5a and consider an individual offered a compound gamble

[14] These models were proposed in response to later evidence (see Section 5) which suggests behavior is more complex than pure fanning-out theories imply.
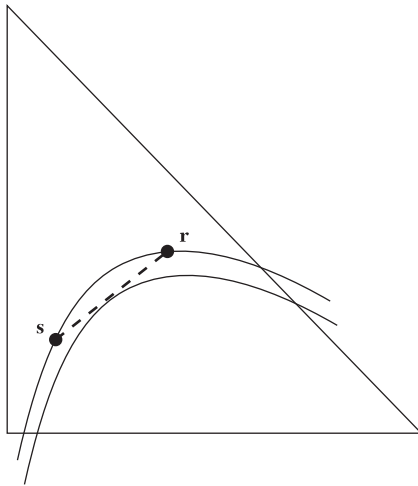
Figure 5b.   Quasi-convex Preferences. Aversion to Randomization

Figure 5c. Quasi-concave Preferences. Preference for Randomization

giving a *p* chance of prospect **q** and a $1 - p$ chance of **r**. Geometrically, the simple prospect induced by this compound gamble must lie along the straight line joining **q** and **r** (for any $0 \leq p \leq 1$). For an individual with linear indifference curves, it follows that for any **q** ~ **r**, the indifference curve through **q** and **r** coincides with the set of simple prospects induced by (**q**, *p*; **r**, $1 - p$). Hence, with linear indifference curves, the individual indifferent between **q** and **r** is also indifferent to randomization between them. Once betweenness is relaxed, this indifference to randomization no longer holds and two important cases can be distinguished: quasi-convex preferences and quasi-concave preferences. A preference function is strictly quasi-convex if for every **q** ≠ **r**, $V(\mathbf{q}, p; \mathbf{r}, (1-p)) <$ max$[V(\mathbf{q}), V(\mathbf{r})]$ for all *p*. When preferences are quasi-convex, indifference curves are *concave*, as in Figure 5b, and consequently the individual will be averse to randomization among equally valued prospects (notice that prospects **r** and **s** in Figure 5b lie on a higher indifference curve than probability mixtures of the two prospects which lie along the dashed line). Conversely, when

preferences are strictly quasi-concave, indifference curves are *convex*, as in Figure 5c, hence by similar reasoning individuals prefer to randomize among equally valued prospects. Some significant theoretical results in economics extend to a non-expected utility world if agents' preferences satisfy betweenness (see Section 5.3 below).

Various models have been proposed that do not impose betweenness. Chew, Larry Epstein, and Segal (1991) propose *quadratic utility theory* which relies on a weakened form of betweenness called *mixture symmetry*: if **q** ~ **r** then (**q**, *p*; **r**, $(1-p)$) ~ (**q**, $(1-p)$; **r**, *p*). In this model, indifference curves may switch from concave to convex (or vice versa) as we move across the triangle. Joao Becker and Rakesh Sarin (1987) propose a model with even weaker restrictions. Their *lottery-dependent utility* assumes only ordering, continuity, and monotonicity. The basic model is conventional theory for minimalists as, without further restriction, it has virtually no empirical content. The authors discuss a particular "exponential form" which implies fanning-out.

An important subset of the betweenness non-conforming theories has an

additional feature absent from the models discussed so far. To this point we have considered a variety of conventional theories, each of which generates the property of fanning-out. Although they achieve it in different ways, there is one structural similarity between these theories: each operates by assigning subjective weights—or utilities—to consequences; the value assigned to any given prospect is then determined by some function that combines these utilities with *objective* probabilities. Another variant of the conventional strategy involves the use of probability transformation functions which convert objective probabilities into subjective *decision weights*. An important feature of these models is that, excepting special cases, betweenness does not hold.

4.1.2 *Theories with Decision Weights*

There is evidence for the view that individuals have subjective attitudes to probabilities which are distinct from attitudes to consequences. For instance, according to Nick Pidgeon et al. (1992), when people are asked to make judgements about the likelihood of death occurring from different causes, they tend to underestimate the number of deaths from relatively frequent causes, while overestimating deaths due to relatively infrequent causes. Similarly, apparent biases in the subjective odds revealed in studies of racetrack betting have been explained as bettors being either oversensitive to the chances of winning on long shots (Mukhtar Ali 1977; Richard Thaler and William Ziemba 1988), or oversensitive to the chances of losing on favorites (Bruno Jullien and Bernard Salanié 1997). These effects might be revealing misperception of objective probabilities or a tendency for individuals to subjectively weight objective probabilities. Either way, in principle,

such effects could be captured in models incorporating *decision weights*. A number of such theories can be understood as variants of the following functional form where the w_i terms represent decision weights:

$$V(\mathbf{q}) = \Sigma_i w_i \cdot u(x_i). \qquad (5)$$

I will call this the *decision weighted form.* Theories of this type were first discussed by Ward Edwards (1955, 1962). In its most basic form, consequences are treated in the way that probabilities are handled in the standard theory and enter "raw" with $u(x_i) = x_i$ for all i. Edwards called this *subjective expected value,* and in the version presented by Jagdish Handa (1977) the decision weight attached to each outcome is determined by a *probability weighting function* $\pi(p_i)$ which transforms the individual probabilities of each consequence directly into weights. As in most theories that incorporate probability weights, $\pi(.)$ is assumed to be increasing with $\pi(1) = 1$ and $\pi(0) = 0$, and I will retain these assumptions from now on. The subjective expected value form has not been widely used, but theories that allow nonlinear transformations of both probabilities and consequences have received much more attention. In the simplest variant of this latter type of model, individuals are assumed to maximize the function:

$$V(\mathbf{q}) = \Sigma\pi(p_i) \cdot u(x_i). \qquad (6)$$

I will call this form *simple decision weighted utility.*[15] Both this and subjective expected value, because they transform the probabilities of individual consequences *directly* into weights (i.e., $w_i = \pi(p_i)$), have the property that $V(\mathbf{q})$ will not generally satisfy monotonicity. To see this, suppose for the sake of example that $\pi(.)$ is convex, then apart from the

[15] This form has sometimes been called *subjective expected utility*, but this label is now more commonly used to refer to Leonard Savages' (1954) formulation of EUT.

extremes of the probability scale, $\pi(p) + \pi(1 - p) < 1$ and there will be some $\varepsilon > 0$ such that gambles of the form $(x, p; x + \varepsilon, 1 - p)$ will be rejected in favor of $(x, 1)$ even though they stochastically dominate the sure option. A similar argument applies for *any* departure from linearity, and the only way to ensure general monotonicity in this type of theory is to set decision weights equal to objective probabilities (i.e., $w_i = \pi(p_i) = p_i$ for all i) in which case the theory reduces to EUT. This property was first noted by Fishburn (1978) and since then has been widely viewed as a fatal objection to models that attach decision weights to the raw probabilities of individual consequences. For example, Machina (1983, p. 97) argues that any such theory will be, "in the author's view at least, unacceptable as a descriptive or analytical model of behavior." The point seems to have been generally accepted and, while many theorists have wished to retain the idea that probabilities may be subjectively weighted, the thrust of work in this stream of the literature over the last two decades has been towards variants of the decision weighted form that satisfy monotonicity.

There are two distinct strands to this contemporary literature: one conventional, the other distinctly nonconventional. The nonconventional route is that taken by Kahneman and Tversky (1979) in *prospect theory*, but that model takes us outside the bounds of conventional theory and so I postpone further discussion of it until the next section. Theorists following the conventional route have proposed decision weighting models with more sophisticated probability transformations designed to ensure monotonicity of $V(.)$. One of the best-known models of this type is *rank-dependent expected utility theory*, which was first proposed by John Quiggin (1982). Machina (1994) describes the rank-dependent model as "the most natural and useful modification of the classical expected utility formula" and, as testament to this, it has certainly proved to be one of the most popular among economists. In this type of model the weight attached to any consequence of a prospect depends not only on the true probability of that consequence but also on its ranking relative to the other outcomes of the prospect. With consequences indexed as before such that x_1 is worst and x_n best, we can state rank-dependent expected utility theory as the hypothesis that agents maximize the decision weighted form with weights for $i = 1,..., n - 1$ given by:

$$w_i = \pi(p_i + \ldots + p_n) - \pi(p_{i+1} + \ldots + p_n)$$

and $\qquad w_i = \pi(p_i) \qquad$ for $i = n$.

In this model there is a meaningful distinction between decision weights (w) and probability weights (π). Richard Gonzalez and George Wu (1999, p. 135) suggest an interpretation of the probability weighting function as reflecting the underlying "psychophysics of risk," that is, the way that individuals subjectively "distort" objective probabilities; the decision weight then determines how the probability weights enter the value function $V(.)$. Notice that $\pi(p_i + \ldots + p_n)$ is a subjective weight attached to the probability of getting a consequence of x_i or better, and $\pi(p_{i+1} + \ldots + p_n)$ is a weight attached to the probability of getting a consequence better than x_i, hence in this theory $\pi(.)$ is a transformation on cumulative probabilities. This procedure for assigning weights ensures that $V(.)$ is monotonic. It also has the appealing property that, in contrast to the simple decision weighting models which assign the same decision weight to any consequence with probability p, the weight attached to a consequence may vary according to how "good" or "bad" it is. So in principle this would allow for, say,

extreme outcomes to receive particularly high (or low) weights. A less appealing feature of the model is that a small change in the value of some outcome of a prospect can have a dramatic effect on its decision weight if the change affects the rank order of the consequence; but a change in the value of an outcome, no matter how large the change, can have no affect on the decision weight if it does not alter its rank.

The predictions of the rank-dependent model depend crucially on the form of $\pi(.)$. If $\pi(.)$ is convex, this generates a set of concave indifference curves (implying aversion to randomization) which are parallel at the hypotenuse but fan out as we move left to right across the triangle and fan in (i.e., become less steep) as we move vertically upwards. Aside from the hypotenuse parallelism which holds for any $\pi(.)$ (see Camerer 1989), the reverse pattern of indifference curves (i.e., convex curves, horizontal fanning in, and vertical fanning out) is generated with a concave $\pi(.)$.

Curvature of $\pi(.)$ in the rank-dependent model has been interpreted as reflecting "optimism" and/or "pessimism" with respect to probabilities (see Quiggin 1982; Menahem Yaari 1987; Enrico Diecidue and Peter Wakker 1999). Consider, for example, the prospect $\mathbf{q} = (x_1, 0.5; x_2, 0.5)$. Assigning weights to the consequences of $\mathbf{q}$ according to the rank-dependent method above gives $w_1 = 1 - \pi(0.5)$ and $w_2 = \pi(0.5)$. With $\pi(.)$ convex, $\pi(0.5) < 0.5$, hence the weight attached to the lower ranking consequence, x_1, will be higher than the weight attached to the larger consequence. This overweighting of the lower-ranked consequences relative to higher-ranked consequences can be interpreted as a form of pessimism. Pessimism also has a close connection to risk aversion: a pessimistic agent with a con-
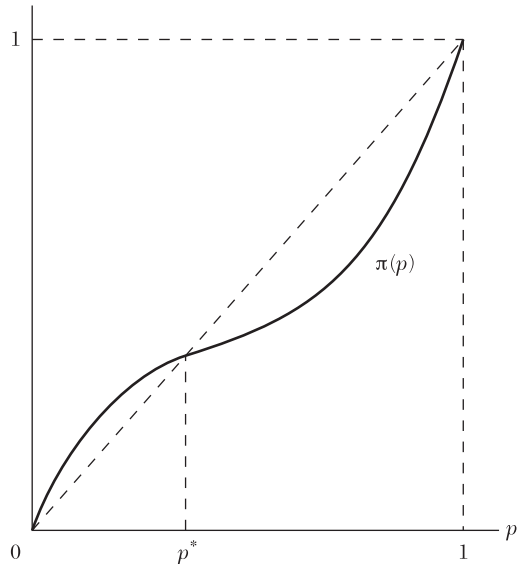
Figure 6. An (inverted) S-shaped Probability Weighting Function

cave $u(.)$ will be universally risk averse; and an agent with a convex utility function can be risk averse if they are sufficiently pessimistic (See Chew, Karni, and Safra 1987; Alain Chateauneuf and Michèlle Cohen 1994).

Although rank-dependent theory does not imply generalized fanning out, the early evidence of EUT violation can be explained either by assuming a simple convex $\pi(.)$ or by more complex specifications. One possibility is the function displayed in Figure 6 which has $\pi(p) = p$ for a unique value of $p = p^*$; it is concave below p^* and convex above it, hence "low" probabilities (below p^*) are overweighted. Quiggin (1982) proposes this form with $p^* = 0.5$. He is drawn to this partly because it explains the early violations of EUT and partly because it has the appealing property that 50–50 bets will be undistorted by probability weighting. While there is little empirical support for the crossover at $p = 0.5$, research over a period of fifty years, from Malcolm Preston and Phillip Baratta (1948) to Drazen Prelec (1998),

lends support to the hypothesis of an (inverted) s-shaped decision-weighting function (see Section 5.1.1). A useful discussion of the theoretical properties necessary and sufficient for an s-shaped weighting function can be found in Tversky and Wakker (1995).

Axiomatizations of rank-dependent expected utility have been presented by, among others, Segal (1990), Wakker (1994), Mohammed Abdellaoui (1999), and Yaari (1987), who examines the special case of the model with linear utility (this is essentially a rank-dependent reformulation of Handa's proposal with $u(x_i) = x_i$). Wakker, Ido Erev, and Elke Weber (1994) provide a useful discussion of the axiomatic foundations of rank-dependent expected utility in which they demonstrate the essential difference between EUT and rank-dependent expected utility is that the latter theory relies on a weakened form of independence called *co-monotonic independence*. It is an implication of the standard independence axiom that if two prospects **q** and **r** have a common outcome x, which occurs with probability p, in each prospect, then substituting x for some other outcome y in both prospects will not affect the preference order of **q** and **r**. The same may not be true in the rank-dependent model, however, because such substitutions may affect the rankings of consequences and hence the decision weights. Co-monotonic independence asserts that preferences between prospects will be unaffected by substitution of common consequences so long as these substitutions have no effect on the rank order of the outcomes in either prospect.

Various generalizations of the rank-dependent model have been proposed (e.g. Segal 1989, 1993; Chew and Epstein 1989; Jerry Green and Jullien 1988). In Green and Jullien, the crucial axiom is ordinal independence. Suppose two prospects **q**, **r** have a "common tail" such that for some j, $p_{qi} = p_{ri}$ for all i from j to n. Ordinal independence requires that preferences between **q** and **r** be unaffected by the substitution of this common tail, in both prospects, with any other common tail. This axiom is necessary for any rank-dependent model. The contribution of Chew and Epstein constructs a theoretical bridge between the rank-dependent models and the betweenness-conforming theories (i.e., those with linear indifference curves discussed above) by presenting a general model which contains each class as a special case (see also the "correction and comment" by Chew, Epstein, and Wakker 1993).

A further extension to the rank-dependent model discussed by Starmer and Sugden (1989), Luce and Fishburn (1991), and Tversky and Kahneman (1992) involves a distinction between consequences that are "gains" and those that are "losses." This approach draws on Kahneman and Tversky's earlier work on prospect theory. It is to this model that we now turn, and in doing so we cross the boundary into nonconventional territory.

4.2 *Nonconventional Theories*

4.2.1 *The Procedural Approach and Reference Dependence*

Each of the theories we have considered so far models choice as preference maximization and assumes that agents behave *as if* optimizing some underlying preference function. The "as if" is significant here: the conventional approach, interpreted descriptively, seeks to predict which choices *are* made and typically, there is no presupposition that the model corresponds with any of the mental activities actually involved in making choices. While this underlying methodology dominates economic theory,

another approach more common in the psychology literature seeks to model the processes that lead to choice. I will call such theories *procedural theories*. A common feature of such theories is to assume that agents draw on decision heuristics or rules of one kind or another when making their choices. The problem is then to identify the set of decision heuristics the agent may draw on, and to specify the conditions under which particular rules will be followed. In such theories, it is common for problem context to be an important determinant of choice-rule selection. For instance, there may be a tendency to choose the rule that is easiest to apply in the given context, and ease of application may depend on how a problem is presented. Consequently, it seems natural to expect phenomena like framing effects within this framework.

One recent and quite general procedural model has been developed by John Payne, James Bettman, and Eric Johnson (1993). They assume that agents have at their disposal a range of possible choice heuristics that might be applied to a given decision task. These include expected utility calculations, satisficing rules, lexicographic choice rules, and so on.[16] In their *adaptive* model the decision maker "decides how to decide," trading off the desire to make a "good" decision against the cognitive effort involved in applying different rules in a given context. Here, as in other procedural models, the agent is conceived of as *boundedly rational*, an agent with limited computational ability and, perhaps, imperfectly defined objectives, attempting to cope with an often complex decision environment. Yet, boundedly rational does not equate with dumb. Payne, Bettman, and

Johnson argue that selection of choice procedures is "adaptive and intelligent" (p. 14), and though decisions may not be optimal in the conventional sense, the selection of decision rule does involve optimization but with unusual constraints (e.g. information processing capacity) and/or objectives (e.g. the choice of strategy might be influenced by considerations such as a desire to be able to justify a choice to a third party). Indeed, as John Conlisk (1996, p. 672) points out, "bounded rationality is not a departure from economic reasoning, but a needed extension of it."

While models of bounded rationality have been applied with some success elsewhere in economics—see Conlisk's (1996) review in this journal—full-blown procedural models of decision under risk, like that of Payne, Bettman, and Johnson, have not received much attention from the economics profession. Nevertheless, there has been a degree of cross-fertilization, and some theories involving a procedural element have appeared in the economics literature. Examples include the models proposed by Kahneman and Tversky (1979), Ariel Rubinstein (1988), and Marc Lavoie (1992).

The most widely discussed of these is Kahneman and Tversky's (1979) *prospect theory*. In this theory, choice is modeled as a two-phase process. In the first phase, prospects are "edited" using a variety of decision heuristics; in the second, choices among edited prospects are determined by a preference function which, for a restrictive class of prospects,[17] can be represented by the

---

[16] For a discussion of satisficing rules see Herbert Simon (1955) and for an example of a lexicographic procedure see Tversky (1969).

[17] The original version of prospect theory does not provide a general preference representation over prospects. Strictly speaking, it only applies to prospects of the form $(x_1, p_1; x_2, p_2; 0, (1 - p_1 - p_2))$. The function assumed in prospect theory coincides with the function defined here in the case of "regular prospects" where either $p_1 + p_2 < 1$, or $x_1 \geq 0 \geq x_2$, or $x_1 \leq 0 \leq x_2$.

Value of Outcome

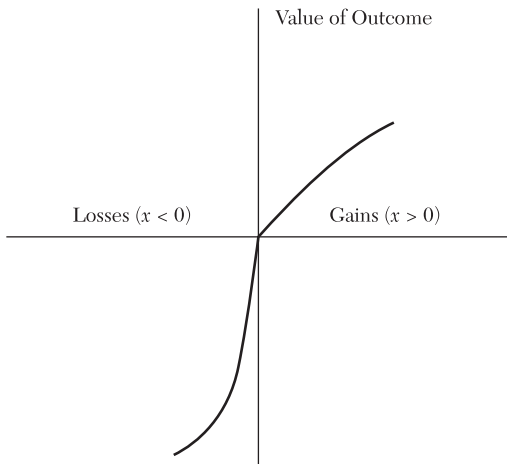Losses (*x* < 0)   Gains (*x* > 0)

Figure 7. The Valuation of Outcomes in Prospect Theory

simple decision-weighted utility form defined in 6 above. Two features of this theory distinguish it clearly from any of the theories we have discussed so far. First and most obvious is the editing phase, but a second distinguishing feature is that, in prospect theory, outcomes are interpreted as gains and losses relative to a *reference point*. For present purposes we may think of the reference point as status quo wealth. The motivation for handling consequences in this way is that it allows gains and losses to be evaluated quite differently. This capacity, it turns out, has some quite interesting implications.

In prospect theory outcomes are evaluated via a utility function[18] with the shape of that in Figure 7. It is kinked at the reference point (i.e., status quo, $x = 0$) and notice two further properties: (i) it is concave for gains and convex for losses, and (ii) it is

---

[18] Kahneman and Tversky explicitly avoid using the term "utility" to describe this function, preferring instead the term "value function." I suspect they had in mind a conception of value independent of risk and wished to distance themselves from the notion of utility in EUT where utilities may partly reflect attitudes to chance. Here I revert to utility terminology, but with a timely reminder that the appropriate interpretation of "utility" varies between theories.

steeper in the domain of losses. In their later paper Tversky and Kahneman (1992) interpret these restrictions as implications of two more general properties of perception and judgement: *diminishing sensitivity* and *loss aversion*. Diminishing sensitivity holds that the psychological impact of a marginal change will decrease as we move further away from a reference point. So, for example, relative to the status quo, the difference between a gain of $10 and $20 will seem larger than the difference between gains of $110 and $120. More generally, the assumption of diminishing sensitivity applied to the outcome domain entails diminishing marginal utility for gains (i.e., $u''(x) \leq 0$ for $x \geq 0$) and diminishing marginal disutility for losses (i.e., $u''(x) \geq 0$ for $x \leq 0$). So property (i) of the utility function is a direct implication of diminishing sensitivity. Loss aversion is the principle that "losses loom larger than corresponding gains" (Tversky and Kahneman 1992, p. 303). They justify this second feature of the function partly by an appeal to intuition and partly by appealing to empirical evidence (e.g. the fact that most people find symmetric bets of the form $(x, 0.5; -x, 0.5)$ "distinctly unattractive"). Loss aversion is modeled by imposing $u'(x) < u'(-x)$.

The evaluation of risky prospects involves a probability weighting function and, in the original version of prospect theory, Kahneman and Tversky proposed a weighting function that underweights "large" and overweights "small" probabilities. The endpoints are such that $\pi(1) = 1$ and $\pi(0) = 0$, but the function is not defined for probabilities close to zero and one; unusual things may happen in these regions—for example, "very small" probabilities might be ignored. It is worth noting that in a later version of prospect theory (see cumulative prospect theory below), Kahneman

and Tversky adopt the widely used inverted-s weighting function. This is partly because that specification fits their data well, and no doubt partly to resolve the ambiguity about what happens at the end points in the original version, but there is also an underlying theoretical rationale. The principle of diminishing sensitivity, which determines some of the important characteristics of the utility function, can also provide a psychological rationale for an (inverted) s-shaped probability weighting function: a function with the property of diminishing sensitivity will be steepest close to a reference point, hence on the assumption that the end points of the probability scale constitute natural reference points, diminishing sensitivity implies a probability weighting function that is steep near zero and one but relatively flat around the middle. The inverted-s has precisely these properties. Hence, if diminishing sensitivity is a general principle of perception, it provides a common psychological underpinning for properties of both the utility function and the probability weighting function.

Kahneman and Tversky (1979) argue that their theory is able to capture a wide range of observed behavior toward risk, including standard violations of the independence axiom (e.g. the common ratio and common consequence effects), and a variety of field data, plus an extensive range of data generated from their own experiments. The theory also has some unusual properties, one of which is the so-called *reflection effect*. The fact that concavity of the utility function in the domain of gains is mirrored by convexity in the domain of losses means behavior towards risk can be likewise mirrored across the two domains. For instance, a given individual who displays risk aversion in a choice among particular prospects with non-

negative outcomes may display risk seeking if all outcomes are changed to losses of the same absolute magnitude. Kahneman and Tversky report evidence for this kind of effect from an experiment involving choices among prospects of the form $s_5 = (x, p; 0, 1 - p)$ and $r_5 = (y, \lambda p; 0, 1 - \lambda p)$. For given absolute values of x and y the majority of subjects revealed $s_5 > r_5$ when $y > x > 0$ and $r_5 > s_5$ when $y < x < 0$.

The "Asian disease" example discussed at the end of Section 3 is consistent with the reflection effect. In that example, the choice between prospects was affected by the description of options. When outcomes were framed as lives saved, the majority of choosers were attracted to a sure gain of 200 out of 600 lives; when framed as losses the majority rejected the sure loss of 400 out of 600 deaths, preferring instead to take the risk. The effect observed there can be interpreted as a reflection effect with risk aversion in relation to gains and risk seeking for losses. Before we could think this an explanation of the Asian disease problem, however, we need an account of how consequences are *interpreted*. From an objective standpoint, two hundred lives saved out of six hundred *is the same thing as* four hundred lives lost, hence a full explanation would require a theory of how framing affects whether an outcome is interpreted as a gain or a loss. Kahneman and Tversky go some way towards this in their discussion of editing.

Prospect theory assumes that prior to the second stage of evaluation, individuals will edit prospects using a variety of heuristics. One of the major editing operations involves the *coding* of outcomes as gains and losses relative to a reference point. Kahneman and Tversky argue that the reference point will typically be the current asset position, but they allow the possibility that "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" (p. 274). Notice that this possibility of differential coding under the two problem descriptions is a necessary step in explaining responses to the Asian disease problem. While some economists might be tempted to think that questions about how reference points are determined sound more like psychological than economic issues, recent research is showing that understanding the role of references points may be an important step in explaining real economic behavior in the field (see, for example, Chip Heath, Steven Huddart, and Mark Lang 1998).

Several of the other editing routines in prospect theory are essentially rules for simplifying prospects and transforming them into a form that can be more easily handled in the second phase. One such operation is the rule of *combination* which simplifies prospects by combining the probabilities associated with identical outcomes. For example, a prospect *described* as $(x_1, p_1; x_1, p_2; x_3, p_3; . . .)$ may be *evaluated* as the simplified prospect $(x_1, (p_1 + p_2); x_3, p_3; . . .)$. Notice that these two prospects are not, in general, equivalent if $\pi(.)$ is nonlinear. Decision makers may also simplify prospects by rounding probabilities and/or outcomes. Further operations apply to *sets* of prospects. The operation of *cancellation* involves the elimination of elements common to the prospects under consideration. Hence a choice between prospects $\mathbf{q'} = (x, p; \mathbf{q}, 1 - p)$ and $\mathbf{r'} = (x, p; \mathbf{r}, 1 - p)$ *may* be evaluated as a choice between $\mathbf{q}$ and $\mathbf{r}$. Although cancellation is effectively an application of the independence axiom of EUT, the editing phase does not imply that choices will generally satisfy independence, since

whether a particular rule is applied depends upon whether or not it is salient. Although they have no formal theory of salience they do present evidence that editing is context dependent. One example shows that cancellation is used in some cases where it is salient and not in others (see their discussion of the "isolation effect," p. 271).

One further rule—I will call it the *dominance heuristic*—has the effect of eliminating stochastically dominated options from the choice set prior to evaluation. The addition of the dominance heuristic does not, however, remove all possibility of monotonicity violation. Kahneman and Tversky assume that individuals scan the set of options and delete dominated prospects *if they are detected*. This ensures the deletion of "transparently" dominated options, but leaves open the possibility that some dominated options survive application of the routine. Since the preference function is not generally monotonic, such options may ultimately be chosen.

This strategy for imposing monotonicity has the further, perhaps surprising, implication that choices may be non-transitive. If $\pi(.)$ is nonlinear, then prospect theory implies that there will be some $\mathbf{q}$ and $\mathbf{r}$ where $\mathbf{q}$ stochastically dominates $\mathbf{r}$ such that $V(\mathbf{r}) > V(\mathbf{q})$.[19] So long as this dominance is transparent, the dominance heuristic ensures that there will be no *direct* violation of monotonicity and $\mathbf{r}$ will not be chosen over $\mathbf{q}$. In general, however, it should be possible to find some other prospect $\mathbf{s}$, such that $V(\mathbf{r}) > V(\mathbf{s}) > V(\mathbf{q})$. If there

---

[19] To see how nonlinearity of $\pi(.)$ can generate violations of monotonicity, consider a simple case where $\mathbf{q} = (x, 1)$ and $\mathbf{r} = (x - \varepsilon, p; x, 1 - p)$. Suppose $\varepsilon > 0$ hence $\mathbf{q}$ dominates $\mathbf{r}$: If $\pi(.)$ is concave, probabilities are overweighted, and the dominated option $\mathbf{r}$ is preferred for some ε. Now suppose $\varepsilon < 0$, hence $\mathbf{r}$ dominates $\mathbf{q}$: if $\pi(.)$ is convex, probabilities are underweighted, and the dominated option $\mathbf{q}$ is preferred for some ε.

is no relation of dominance between **s** and either of **q** or **r**, then pairwise choice among these three gambles will generate a systematic cycle of choice in which **q** $>_c$ **r** and **r** $>_c$ **s** and **s** $>_c$ **q** where $>_c$ is the relation "is chosen over." Quiggin (1982, p. 327) calls this an "undesirable result." Quiggin's reaction would not be untypical of economists more generally, most of whom have taken both transitivity and monotonicity to be fundamental principles which any satisfactory theory should embody. On the other hand, several economists, Quiggin included, have thought aspects of prospect theory appealing and have sought to build the relevant features into models more in keeping with conventional theoretical desiderata. For example, part of Quiggin's motivation in developing rank-dependent expected utility theory was to establish that a central feature of prospect theory—nonlinear decision weights—can be built into a preference function without sacrificing monotonicity. By constructing decision weights cumulatively, we obtain a (transitive) preference function that is monotonic without the need for an additional editing routine. Papers by Starmer and Sugden (1989), Luce and Fishburn (1991), and Tversky and Kahneman (1992) show that the rank-dependent form can be extended to capture another key element of prospect theory: valuing outcomes relative to *reference points*.

In Starmer and Sugden's model, any prospect **q** is valued by the function $V(\mathbf{q}) \equiv V^+(\mathbf{q}) + V^-(\mathbf{q})$ where $V^+(\mathbf{q})$ is the rank-dependent expected utility of a *transformed* prospect $\mathbf{q}^+$; this is equivalent to **q** excepting that any outcomes of **q** that are *losses* are replaced by zeros. Similarly, $V^-(\mathbf{q})$ is obtained by applying the standard rank-dependent form to a transformed prospect $\mathbf{q}^-$; in this case, any outcomes that are *gains* are re-

placed by zeros. Tversky and Kahneman's model, *cumulative prospect theory*, is more general in that it allows the decision weighting function to be different for the positive and negative components. The development of these so-called sign- and rank-dependent models demonstrates that important aspects of prospect theory can be captured within a formal model that is essentially conventional, without the need to invoke an editing phase.

In these later models, the procedural element central to prospect theory has disappeared.[20] No doubt the abandonment of editing does leave some things unexplained. For instance, framing effects do suggest that choices are context dependent in complex yet subtle ways, and the procedural approach seems to provide the more natural arena in which to model this. On the other hand, introducing elements of bounded rationality does considerably complicate the theoretical structure of models in ways that render them less compatible with the rest of economic theory. For example, working with a set of decision rules seems clumsy, relative to the neatness and tractability of optimizing a single function; unlike conventional models, procedural models often exhibit a degree of indeterminacy.[21]

Might such arguments provide sufficient grounds for defending a general theoretical presumption that agents behave "as if" fully rational? Conlisk (1996) reviews a series of methodological arguments which might be used to make such a case against incorporating

----

[20] Although Tversky and Kahneman do mention that editing may be important, their 1992 model has no formal editing phase and their references to it are virtually asides.

[21] For instance, in prospect theory, the outcome of editing can depend on factors that are underdetermined by the theory, such as the order in which operations are applied (see M. K. Stevenson, J. R. Busmeyer, and J. C. Naylor 1991).

ideas of bounded rationality into economics. He concludes that it is hard to make any convincing case against. If that's correct, and I for one am persuaded, then the question to ask is whether departures from conventional models are of sufficient concern, from an empirical point of view, to justify the theoretical costs involved. We shall be examining this issue shortly (see Section 5). First, however, we consider an alternative avenue of departure from the conventional approach.

4.2.2 *Non-Transitive Preference Theory*

As we have seen, many have taken the view that the standard independence axiom of EUT can be sacrificed for the sake of explaining the data. Transitivity, however, may be another matter. It might be tempting to think that transitivity is so fundamental to our ideas about preference that to give it up is to depart from theories of preference altogether. Can we speak of people maximizing anything if they don't have transitive preferences? It turns out that the answer is yes.

There is at least one well-known theory of choice based on a model of non-transitive preference. The theory I have in mind was proposed simultaneously by Bell (1982), Fishburn (1982), and Loomes and Sugden (1982). I will begin by discussing a version of this theory presented by Loomes and Sugden (1987). Loomes and Sugden call their theory *regret theory,* and its central premise is closely akin to the psychological intuition at the heart of the theory of disappointment. In that theory, it is assumed that an individual compares the outcomes *within a given prospect* giving rise to the possibility of disappointment when the outcome of a gamble compares unfavorably with what they might have had. Regret theory allows comparisons between consequences to affect choice, but in this case, the relevant compari-

sons occur between the consequences of alternative choice options.

Since the theory has to allow comparisons between choice options, it cannot be a conventional theory that assigns values independently to individual prospects. Loomes and Sugden propose a theory of pairwise choice in which preferences are defined over pairs of *acts*, where an act maps from *states of the world* to consequences.[22] Let Ai and Aj be two potential acts that result in outcomes x_{is} and x_{js}, respectively, in state of the world S. The utility of consequence x_{is} is given by a function $M(x_{is},x_{js})$ which is increasing in its first argument and decreasing in its second. This function allows the utility from having x_{is} to be suppressed by "regret" when $x_{is} < x_{js}$, or enhanced by "rejoicing" when $x_{is} > x_{js}$. The individual then seeks to maximize $\Sigma_s\, p_s\, .M(x_{is},x_{js})$ where p_s is the probability of state S. Regret theory reduces to EUT in the special case where $M(x_{is},x_{js}) = u(x_{is})$.

Although preferences are defined over acts, the theory can be applied to choices between prospects given some assumption about how outcomes are correlated between them. One interesting case is when consequences are uncorrelated between prospects; that is, when prospects are *statistically independent*. In a choice between a pair of such prospects **q** and **r**, *if* **q** is chosen, the probability of getting x_i and missing out on x_j is given by $p_{qi}p_{rj}$ where p_{qi} is the probability of consequence x_i in **q** and p_{rj} the probability of x_j in prospect **r**. Preferences between **q** and **r** are then determined by the expression:

$$\mathbf{q} \underset{<}{\overset{>}{\sim}} \mathbf{r} \Leftrightarrow \Sigma_i\Sigma_j p_{qi}p_{rj}\psi(x_i, x_j) \underset{<}{\overset{>}{=}} 0 \qquad (7)$$

---

[22] As a theory of pairwise choice, regret theory has limited applicability, but ways of generalizing the theory have been suggested by Sugden (1993) and Quiggin (1994).

where $\psi(x_{is}, x_{js}) \equiv M(x_{is}, x_{js}) - M(x_{js}, x_{is})$. The function $\psi(.,.)$ is skew symmetric by construction, hence $\psi(x,y) \equiv -\psi(y,x)$ and $\psi(x,x) \equiv 0$ for all x,y.

If prospects are statistically independent, the addition of a further assumption which Loomes and Sugden call *regret aversion*[23] implies that indifference curves will fan out in the probability triangle. Regret aversion requires that for any three consequences $x > y > z$, $\psi(x, z) > \psi(x, y) + \psi(y, z)$. The interpretation of the assumption is that large differences between what you get from a chosen action and what you might have gotten from an alternative give rise to disproportionately large regrets; so people prefer greater certainty in the distribution of regret. Under these conditions regret theory is equivalent to Chew and MacCrimmon's weighted utility theory, and so indifference curves in the probability triangle will have the pattern described in Figure 4 above (see Sugden 1986 for a simple demonstration of this). Consequently, regret theory is able to explain the standard violations of the independence axiom for statistically independent prospects.[24]

If we consider the class of all statistically independent prospects—not just those with up to three pure consequences—weighted utility theory is a special case of regret theory. Specifically, the representation in expression 7 is obtained from Chew and MacCrimmon's axiom set by relaxing transitivity. This is the route by which Fishburn (1982) arrived at this model (he calls it

skew-symmetric bilinear utility or SSB). Fishburn's model is identical with regret theory for statistically independent prospects, and we can think of regret theory as a generalization of SSB which extends it to non-independent prospects: in this realm, regret aversion has some very interesting implications.

Consider three stochastically equivalent actions A_1, A_2 and A_3, each of which gives each of the consequences $x > y > z$ in one of three equally probable states of the world s_1, s_2 and s_3. Any conventional theory entails a property of *equivalence*, that is, indifference between stochastically equivalent options, hence, for any such theory, $A_1 \sim A_2 \sim A_3$. In regret theory, however, it matters how consequences are assigned to states, and for particular assignments, regret theory implies a strict preference between stochastically equivalent acts, violating equivalence. For example, suppose that the three acts involved the following assignment of consequences to states:

	s_1	s_2	s_3
A_1	z	y	x
A_2	x	z	y
A_3	y	x	z

If we consider preferences between the first two acts, regret theory implies:

$$A_1 \underset{<}{\overset{>}{\sim}} A_2 \Leftrightarrow [\psi(z,x) + \psi(y,z) + \psi(x,y)] \underset{<}{\overset{>}{=}} 0 \quad (8)$$

Using the skew symmetry of $\psi(.,.)$, the term in square brackets is equal to $[\psi(x,y) + \psi(y,z) - \psi(x,z)]$. Assuming regret aversion, this will be negative, hence regret theory implies a strict preference $A_2 > A_1$. It is easy to see that the same reasoning applied to the other two possible pairwise comparisons implies $A_3 > A_2$ and $A_1 > A_3$. Hence, regret theory also implies a *cycle of preference* of the form: $A_2 > A_1$, $A_3 > A_2$, $A_1 > A_3$. Now consider adding some small positive amount ε to one consequence of action

[23] In their early discussions of regret theory, Loomes and Sugden called this assumption "convexity."

[24] Some instances of the common consequence effect have involved statistically non-independent options, and these cases are not consistent with regret theory (unless we assume agents treat options as if they are independent even when they are not).

A_1. The resulting action, call it A_1^*, stochastically dominates each of the original actions. But since regret theory implies $A_2 > A_1$ we should expect $A_2 > A_1^*$ for at least some $\varepsilon > 0$. Hence regret theory also implies violations of monotonicity.

Relative to the conventional approach then, preferences in regret theory are not at all well-behaved: they satisfy neither monotonicity nor transitivity and the theory allows strict preferences between stochastically equivalent acts. While such properties may seem peculiar to the eye of the conventional economist, from the descriptive angle, the crucial question is whether such implications of the theory are borne out by actual behavior. Shortly after proposing regret theory, Loomes and Sugden (1983) argued that at least one might be. Consider the following three acts labeled $, P and M with monetary consequences $x > y > m > 0$ defined (for the sake of simplicity) over three equiprobable states:

	s_1	s_2	s_3
$	x	0	0
P	y	y	0
M	m	m	m

The actions labeled $ and P have the structure of typical $- and P-bets: they are binary gambles where $ has the higher prize, and P the higher probability of "winning"; the third act gives payoff m for sure. Loomes and Sugden show that, given regret aversion, pairwise choices over acts with this structure may be cyclical, and if a cycle occurs it will be in a specific direction with $P > \$$, $M > P$ and $\$ > M$. Now recall that in a standard experiment, subjects reveal $P > \$$ in a straight choice between options but place a higher value on $ relative to P in separate valuation tasks. If we interpret choices from {$, M} and {P, M} as analogues of valuation tasks asking "is $ (or P) worth more or less than m" then the cycle predicted by regret

theory can be interpreted as a form of preference reversal.

So, regret theory offers the tantalizing opportunity of explaining violations of independence and preference reversal within a theory of preference maximization. Of course, since observation of preference reversal pre-dates the development of regret theory, that phenomenon offers only weak support for the unconventional predictive content of regret theory. More recent research has aimed at testing some novel predictions of regret theory, and some of the results from this line of research are discussed in the next section.

5. *Evaluating Alternatives to Expected Utility Theory*

5.1 *The Recent Experimental Evidence*

Starting in the mid-1980s, a number of researchers turned their attention towards testing non-expected utility theories. The majority of this work involved experimental testing, some of it designed to compare the predictive abilities of competing theories; some designed to test novel implications of particular theories; and some designed to test the descriptive validity of particular axioms. A very large volume of work has emerged in this arena, providing a much richer evidential base against which theories can be judged. The purpose of this section is to discuss what has been learned in this second phase of the hunt, and its organization reflects the dichotomy between conventional and nonconventional models adopted in the theoretical discussion. The first part discusses evidence that bears directly on the choice between conventional models; the second part addresses evidence with a bearing on the relative merits of conventional and nonconventional approaches.

5.1.1 *Choosing among Conventional Models*

As we have seen, conventional theories all imply the existence of indifference curves in the probability triangle, and certain of their key properties can be expressed in terms of characteristics of the indifference maps they generate. For instance, Machina's theory implies generalized fanning-out, while other theories imply a mixture of fanning-in and fanning-out. A large number of experimental studies have explicitly examined individual behavior in choices among prospects in probability triangles. The data generated from these "triangle experiments" provide a vantage point from which we can ask the following question: suppose one were attempting to construct a conventional theory now, with the aim of accounting for the evidence currently available, are there any obvious properties one should seek to build in?

Although the evidence is both rich and complex, a number of stylized facts apply across a range of studies. In my view, three observations seem particularly robust. First, if you want a theory consistent with the available data *don't impose generalized fanning-out*. Evidence from a wide range of studies reveals behavior inconsistent with linear parallel indifference curves, but the patterns actually observed are more complex than generalized fanning-out. For example, while numerous studies reproduce behavior consistent with Allais paradox violations of EUT in choice pairs moving left to right along the bottom edge of the probability triangle, another finding replicated across a range of studies, including Camerer (1989), Chew and William Waller (1986), Battalio, Kagel, and Jiranyakul (1990), and Starmer (1992), is a tendency for behavior to become less risk averse moving up along the left-hand edge of probability triangles. Such behavior would be consistent with a tendency for indifference curves to fan in. These facts mitigate in favor of theories like disappointment aversion, implicit utility, quadratic utility, and models with decision weights, all of which allow a mixture of fanning-in and fanning-out.

A second general lesson in the data seems to be *don't impose betweenness.* There is considerable evidence—a good part of it is reviewed in Camerer and Teck-Hua Ho (1994)—that choices are inconsistent with the assumption of linear indifference curves. Together these two requirements narrow the field considerably: if we want a theory of mixed fanning with nonlinear indifference curves, of the theories reviewed above the only contenders are quadratic utility, lottery-dependent utility, and models with decision weights.

A third widely observed finding arguably nudges the decision weighting models into the lead: *behavior on the interior of the probability triangle tends to conform more closely to the implications of EUT than behavior at the borders*. Although significant off-border violations are observed in at least some experiments (see for example Wu and Gonzalez 1996) several studies, including those of Conlisk (1989), Camerer (1992), David Harless (1992), and Garry Gigliotti and Barry Sopher (1993), suggest that violations of EUT are *concentrated* in comparisons between options involving prospects on or near to the borders of triangles. It is important to note that this observation is unlikely to rescue EUT for practical purposes. A natural interpretation of the "border effect" is that individuals are particularly sensitive to changes in the likelihood of outcomes with "extreme" probabilities (i.e., moving off the border of the triangle we introduce a low probability event; in the

vicinity of each corner, some outcome is near certain). It is very easy to think of important choice scenarios involving real prospects with "extreme" probabilities; for example, individual decisions about participation in national or state lotteries or collective decisions about nuclear power generation involve high magnitude outcomes (winning the lottery, suffering the effects of a radiation leak) occurring with very small probabilities. Consequently, there are good reasons to model sensitivity to "extreme" probabilities. One obvious way to do it is via decision weights.[25]

In summary, if one is looking to organize the data from the large number of triangle experiments, then the decision-weighting models are probably the best bet. Moreover, there is a striking degree of convergence across studies regarding the functional form to use; for best predictions the key ingredient seems to be an inverted s-shaped weighting function. Empirical support for this specification comes from a wide range of studies including Pamela Lattimore, Joanna Baker, and Ann Witte (1992); Tversky and Kahneman (1992); Camerer and Ho (1994); Abdellaoui (1998); and Gonzalez and Wu (1999), all of which fit the decision-weighting model to experimental data. Collectively, these studies show that models with s-shaped probability transformations offer significant predictive improvement over EUT and outperform other rivals. Most of the studies in this vein, at least those conducted in recent times, employ the rank dependent transformation method, though different mathematical forms have been used

for the probability weighting function. Lattimore, Baker, and Witte (1992) use a probability weighting function of the form:

$$\pi(p_i) = \alpha p_i^\beta / \left[\alpha p_i^\beta + \sum_{k=1}^{n} p_k^\beta \right]$$

for i, $k = 1,2, \ldots n$, $k \neq i$ and α, $\beta > 0$ (n is the number of outcomes as usual). This captures a number of other proposed forms (e.g. those of Uday Karmarkar 1978 and Quiggin 1982) as special cases. With $\alpha = \beta = 1$, $\pi(p_i) = p_i$, hence we get EUT. More generally, the parameter β controls the inflection point and $\beta < 1$ generates the inverted-s with the consequent over-weighting of "small" probabilities below the inflection point, and underweighting above it. With $\alpha < 1$, $\pi(.)$ is "sub-certain" in the sense that the sum of weights ($\Sigma_i \pi(p_i)$) will be less than unity. Lattimore, Baker, and Witte (1992, p. 381) describe this as "'prospect pessimism' in the sense that the value of the prospect is reduced vis-à-vis certain outcomes." In their empirical estimates, they find that allowing nonlinear decision weights offers significant improvement in predictive power over EUT (which is the best model for only about 20 percent of their subjects). The best-fitting weighting function is generally the inverted-s exhibiting greater sensitivity to high and low probabilities relative to mid-range probabilities. They also report differences between the best-fitting weighting functions for gains and losses (for example "pessimism" is more pronounced for losses) though the interpretation of these differences is potentially confounded by the fact that, in their study, gains are measured in units of money while losses are measured in units of time.

Single parameter weighting functions have been proposed by Tversky and Kahneman (1992) and Prelec (1998). Tversky and Kahneman suggest the

[25] Another theoretical possibility suggested by Neilson (1992) is to allow the utility function defined over outcomes to depend on the number of outcomes: this generates different behavior on and off the border, but experimental tests of the model (see Stephen Humphrey 1998) have not been supportive.

form $\pi(p) = p^\delta/[(p^\delta + (1 - p)^\delta)^{1/\delta}]$. This generates the inverted-s for $0 < \delta < 1$, and reducing δ lowers the crossover point while accentuating the curvature of the function. Their empirical analysis supports the s-shaped weighting function and also reveals systematic differences in behavior for gains and losses; specifically, indifference curves in the best-fitting models for losses resemble those for gains flipped around a 45-degree line. This supports the case for a model that distinguishes between gains and losses (i.e. a model with a reference point) though virtually no work is done by the weighting function here; essentially, the same probability-weighting function works well for both gains and losses.

Prelec proposes the function $\pi(p) = \exp(-(-\ln p)^\alpha)$. With $0 < \alpha < 1$, this generates the inverted-s with a fixed inflection point at $p = 1/e = 0.37$. Visually, α is the slope of $\pi(.)$ at the inflection point, and as α approaches unity, $\pi(.)$ becomes approximately linear; as it approaches zero, $\pi(.)$ approximates a step function. Prelec argues that a crossover in the vicinity of $1/e$ is consistent with the data observed across a range of studies. A novel feature of Prelec's contribution is to provide an axiomatization for this form, and he also discusses a two-parameter generalization. The two-parameter version is similar in spirit to the "linear in log odds form" discussed by Gonzalez and Wu (1999) in that it allows the curvature and elevation of the weighting function to be manipulated (more or less) independently. In the latter form, probability weights are given by:

$$\pi(p_i) = \delta p_i^\gamma/[\delta p_i^\gamma + (1 - p_i)^\gamma]$$

The parameter δ primarily controls the absolute value of $\pi(.)$ by altering the elevation of the function, relative to the 45-degree line, while γ primarily controls curvature. Gonzalez and Wu's data suggests that the flexibility of a two-pa-

rameter model may be useful for explaining differences between individuals. For other purposes, however, parsimony favors the one-parameter versions.

Conventional theory can claim a success here: a one-parameter extension to EUT can offer significantly improved predictive power for a large body of data generated mainly from triangle experiments. If we want to predict behavior over simple choices like this we know a lot about how to improve on EUT. Reflection over a broader range of experimental evidence, however, suggests that we are still a long way from a satisfactory general account of behavior under risk.

5.1.2 *A Case for the Unconventional*

I now turn to a discussion of laboratory evidence which, in my view, provides a substantive challenge to certain key assumptions that underpin conventional preference theories. Again, the evidence I cite does not constitute a thorough review; instead I draw on examples of phenomena that seem both challenging and well established.

i. Violations of monotonicity: It might be tempting to think that violations of monotonicity must be rare for two reasons based on casual empiricism: individuals are not stupid; and, if they were, we would see market institutions trading on that stupidity (e.g. casinos competing with each other by advertising worse odds of winning than their rivals!). Experimental evidence, however, supports two stylized facts about monotonicity. First, *very few people will choose a stochastically dominated option from a choice set when it is transparently obvious that the option is dominated*. Second, *choices are not generally monotonic* and systematic violations of monotonicity can be generated in contexts where the relation of dominance is opaque (i.e. not obvious to the

chooser). One sharp illustration of this is provided by the following example due to Tversky and Kahneman (1986):

Consider the following pair of 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 very easy to see that option B dominates option A since, for every color, the prize for option B is always at least as good as the prize for option A and in some cases it is better. Kahneman and Tversky presented this problem to 88 subjects and found that all of them chose B. Now consider this slightly modified version of the above problems:

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

Options C and D are stochastically equivalent to A and B respectively; the only difference is a minor change in the presentation which "simplifies" the options by assigning each prize to a single color. This framing of the options, however, also makes it more difficult to detect the dominance of D over C. In fact, Kahneman and Tversky found that a majority of subjects (58 percent) chose the dominated option C. This finding—which is consistent with the original two-phase version of prospect theory—supports the view that although people do not purposefully choose to violate monotonicity, they might do so in cases where

the violation is opaque, presumably because they do not have generally monotonic preferences. Further examples of monotonicity violation can be found in Loomes, Starmer, and Sugden (1992); Michael Birnbaum and Laura Thompson (1996); Birnbaum and Juan Navarette (1998); and J. W. Leland (1998).

ii. Event-splitting effects: It is well known in marketing circles that by "unpacking" positive attributes of a good into multiple sub-attributes one can make a good seem more desirable. For example, instead of just describing a car as having "good performance," you can make it seem more attractive by subdividing performance into acceleration, cornering, braking, and so on. It turns out that the attractiveness of risky options can be influenced by "unpacking" probabilities in an analogous way. Consider acts defined over a set of mutually exclusive and exhaustive states of the world $S = \{s_i: i = 1, \ldots, n\}$. Define four events E_1, E_2, E_3 and E_4 such that E_1 consists of the set of states $\{s_i: i = 1, \ldots, k - 1\}$; E_2 consists of the remaining states $\{s_i: i = k, \ldots, n\}$. Events E_3 and E_4 partition E_2 into two distinct subsets of states such that $E_3 = \{s_i: i = k, \ldots k + j\}$ and $E_4 = \{s_i: i = k + j + 1, \ldots, n\}$. Now consider two particular acts A and B where A gives consequence x conditional on E_1 and consequence y conditional on E_2; B gives x conditional on E_1 and y conditional on each of events E_3 or E_4. The only difference between A and B is that A is *described* as resulting in outcome y for a single event whereas act B is *described* as resulting in y for two distinct events. In most theories this difference is irrelevant and the two acts are simply regarded as two alternative descriptions of the same prospect $\mathbf{q} = (x, p; y, 1 - p)$ where p is the probability of E_1.

There is clear evidence that such re-descriptions do matter, however. More

specifically, studies by Starmer and Sugden (1993) and Humphrey (1995) have shown that when an event that gives some outcome *y* is split into two sub-events, there is a tendency for that consequence to carry more weight even though its total objective probability is unchanged. This is the event-splitting effect. One implication of it is that individuals are more likely to choose particular options if events containing relatively attractive consequences are subdivided into two (or perhaps more) events. Similar effects have been found by Martin Weber, Franz Eisenführ, and Detlof vonWinterfeldt (1988) in the context of multi-attribute choice. This suggests the possibility, albeit speculative, that event splitting may be a consequence of some more general property of judgement which extends beyond the domain of risk. The existence of an event-splitting effect has some potentially important implications at the level of policy. Take, for example, the risks of skin cancer from exposure to ultraviolet radiation (UVR). There are a variety of cancers that may result from UVR. If event-splitting effects generalize to this context, then the perception of the risks from UVR may be affected by whether the risks are described collectively as a single risk of "skin cancer" with a given probability, or alternatively described as a series of risks of different cancers, each occurring with a smaller probability. The evidence from the laboratory suggests that the risks would appear worse under the second, disaggregated description.

One possible explanation for the event-splitting effect is the *support theory* of Tversky and Koehler (1994). This theory distinguishes between events and "hypotheses" which are descriptions of events. "Support" is a measure of the strength of evidence in favor of a hypothesis, and it may be grounded in objective data or subjective impression. The theory explicitly allows for judged probability to be enhanced by "unpacking" of events in line with the event-splitting effect. *Prospective reference theory* proposed by W. Kip Viscusi (1989) provides another possible account of event splitting; this works by assuming that individuals may distrust, and hence revise, objective probability information. Another possible explanation is provided by simple decision-weighted utility: take the formulation in (6) above, then assume a sub-additive weighting function such that $\pi(p_1) + \pi(p_2) > \pi(p_1 + p_2)$ for at least some p_1, p_2. A theory of this form would be consistent with event splitting and also the violations of monotonicity observed by Loomes, Starmer, and Sugden (1992).[26] Historically, of course, theories of this type with such a simple construction of decision weights have been rejected in favor of more sophisticated theories like rank dependent models precisely because they predict such behavior.

iii. Violations of transitivity: There is well-established evidence—Tversky (1969) produced some of the earliest—that cyclical choice is a robust and reasonably general phenomenon. This violates a central principle of most economic theory. Some economists might be tempted to argue that economics can explain non-transitive behavior under risk using regret theory: regret theory is, after all, a model of preference maximization (i.e. in the tradition of economic theory) and it allows non-transitive behavior. The truth is, however, that even regret theory cannot actually explain what we are learning about intransitivities in risky choice.

[26] It has turned out that a large part of the experimental evidence which seemed to support novel predictions of regret theory is more likely due to event-splitting effects. This possibility was first suggested by Stephen Davies during a seminar at the University of East Anglia.

While it is true that experimental studies have found particular forms of intransitivity specifically predicted by regret theory (see Loomes, Starmer, and Sugden 1989, 1991),[27] later experiments by Starmer and Sugden (1998) confirm the robustness of the phenomenon, but also allow us to reject a variety of possible explanations for the occurrence of the cycle, *including the explanation offered by regret theory*. The bottom line is that economists do not have a theory of non-transitive behavior that is consistent with the available evidence, though some of the evidence is suggestive of the kind of theory that would be needed. Here's an example.

Earlier in this paper, I noted that, because the original version of prospect theory combines simple nonlinear decision weights with a dominance heuristic, the theory implies violations of transitivity. Some have thought that a limitation of the theory, but such judgements may have been premature (and confusing normative and descriptive issues). Starmer (1999a) reports an experiment that tests for the specific form of intransitivity implied by prospect theory and finds it. I would suggest there is a general lesson here that runs beyond simply observing a new form of intransitive behavior: don't judge the predictions of descriptive theories using normative principles of choice; judge them against empirical evidence.

No theory currently available in the economics literature can successfully organize the observations in i–iii above. One might argue it is still early days and that the way forward is to develop a further generation of theories in the light of the accumulating data. It is surely important to acknowledge, however, that this type of evidence is not

just at odds with most available theories. Arguably, it strikes deeper since it might be read as suggesting a flaw in the conventional modeling strategy. In seeking models of actual behavior, theorists following the conventional strategy have sought theories built upon consistency principles like transitivity and monotonicity. Not only has this been the standard approach, but those relatively rare theories that have not conformed with these principles have been widely criticized as unacceptable or implausible. For example, the original version of prospect theory might explain at least some of the above evidence, but these explanations would all rely on elements of prospect theory—like procedural rules, or "unsophisticated" probability weighting functions—which have been criticized, ignored, or abandoned. But, like it or not, it seems that theories of well-behaved preferences in the conventional mould will not provide general descriptive models consistent with the experimental evidence. Consequently, I would argue that if we genuinely seek descriptive models capable of explaining the patterns observed in laboratory behavior, our conventional theoretical desiderata may need rethinking: in particular, there should be no prior supposition that the best models will be ones based on principles of rational choice, no matter how appealing those may seem from a normative point of view.

5.2 *Evidence from the Field*

I have heard some economists argue that they would take more notice of non-EU models if they could be shown cases where they help to explain real-world phenomena of practical interest to economics. It is a fair point, but proponents of non-expected utility theory can muster some strong responses. Let me illustrate by way of a couple of

---

[27] For further evidence and discussion see Loomes and Taylor (1992) and Tversky, Slovic, and Kahneman (1990).

examples. It is well-known that EUT does a poor job explaining some of the things economists are traditionally very interested in, like insurance behavior and the demand for assets: in both cases non-expected utility models may offer a better understanding of the determinants of some real market behavior, though in the second case a nonconventional approach might hold the solution.

The standard theory of insurance based on EUT has some implications that have long been regarded as highly implausible. For example, a risk-averse expected utility maximizer will not buy full insurance in the presence of positive marginal loading (see J. Mossin 1968). This implication, Karl Borch (1974) suggests, is "against all observation." More recently, Wakker, Thaler, and Tversky (1997) have made a similar point in relation to "probabilistic insurance." Think of probabilistic insurance as a policy with some fixed probability q that a claim will not be paid in the event of an insured loss. Wakker, Thaler, and Tversky show that an expected utility maximizer willing to pay a premium c for full insurance against some risk should be willing to pay a premium approximately equal to the actuarially adjusted premium $(1-q).c$ for probabilistic insurance. Survey evidence, however, shows that people are extremely averse to probabilistic insurance and their willingness to pay for it is much less than standard theory allows.

If expected utility can't explain insurance behavior, can non-expected utility theory do any better? Part of the answer is provided by Segal and Avia Spivak (1990), who show that a number of implications of EUT for insurance and asset demand which are widely recognized to be counter intuitive have a common origin. They arise because, with any smooth (i.e. differentiable) utility function, EUT implies that agents will be approximately risk neutral for small risks (since the utility function will be almost linear). This theoretical property is at odds with peoples' actual risk attitudes as revealed through their reactions to probabilistic insurance and so on: people demand a much greater reduction in premium than the actuarially fair adjustment for accepting a small positive risk of claim nonpayment.

Segal and Spivak go on to show that the counter-intuitive implications of EUT carry through to non-expected utility theories which have similar smoothness properties. This captures a large number of alternatives to EUT and, in fact, only a single type of theory escapes their net: the decision weighting models. It is easy to see why models with probability transformations do not imply approximate risk neutrality for small risks since risk averse behavior can be generated by nonlinear probability weighting even where the utility function is linear. So, for example, aversion to probabilistic insurance is easily explained by overweighting of the small probability of non-payment. As such, decision weighting models stand out as leading contenders to explain aspects of insurance behavior which it has long been known standard theory cannot handle. There is growing evidence that probability weighting may be an important ingredient in explaining a variety of field data relating to gambling and insurance behavior, and several examples are discussed by Camerer (forthcoming).

Another field-phenomenon that has perplexed economists is the size and persistence of the excess return on stocks over fixed income securities. This is the so-called *equity premium puzzle* and it is the economics equivalent of the crop circle: we have seen it in the field, but we have real trouble explaining how it got there. Since the return

on stocks is more variable, standard theory is consistent with some difference in the long-run rates of return, but since Rajnish Mehra and Edward Prescott (1985) it has been recognized that the observed disparity implies implausibly high degrees of risk aversion in standard models of asset pricing. One possible explanation for (part of) the equity premium has been suggested by Epstein and Stanley Zin (1990). They show that a recursive utility model using rank dependent preferences predicts an equity premium, though only about one third of the size that is usually observed. A full—and in my view much more convincing—account has been suggested by Shlomo Benartzi and Thaler (1995), who show that the level of equity premium is consistent with prospect theory, with the added assumption that agents are myopic (i.e., they assess expected returns over "short" time horizons). The crucial element of prospect theory for this explanation is loss aversion. In the short run, there is a significant chance that the return to stocks is negative, so if, as loss aversion implies, investors are particularly sensitive to these possible negative returns, that would explain the equity premium for myopic investors. But just how loss averse and how myopic do agents have to be for this explanation to work? Benartzi and Thaler show that, assuming people are roughly twice as sensitive to small losses as to corresponding gains (which is broadly in line with experimental data relating to loss aversion), the observed equity premium is consistent with the hypothesis that investments are evaluated annually. This is a very simple, and to my mind, intuitively appealing account of another important field phenomenon which has defied explanation in standard theory.

Notice that while loss aversion can be accommodated in conventional models like the sign and rank dependent theories, the other ingredient in this explanation of the equity premium—i.e., myopia—belongs in another tradition. This is essentially a bounded rationality assumption, and while the one-year time horizon has a nice ring of plausibility to it, it sits much more naturally alongside procedural theories like the original version of prospect theory. Bounded rationality assumptions seem to be providing the missing links necessary to explain an increasing range of economic phenomena (see Camerer 1998 for a recent review of applications in individual decision making).

It also seems likely that the concept of loss aversion will become increasingly important in economics. Evidence for the existence of loss aversion in both risky and riskless environments now seems overwhelming—a few of many possible references are B. J. McNeil et al. (1982); Jack Knetsch and J. A. Sinden (1984); William Samuelson and Richard Zeckhauser (1988); Knetsch (1989); George Loewenstein and Daniel Adler (1995); Kaisa Herne (1998)—and loss aversion may well explain other puzzles in field data such as the disparity between measurements of willingness to pay and willingness to accept (see Ian Bateman et al. 1997) plus a variety of other examples relating to consumption and labor supply decisions discussed in Camerer (forthcoming). Yet more examples relating to the evaluation of opportunity costs, sunk costs, and search behavior are discussed in the much earlier paper by Thaler (1980).

These examples show that there are important field phenomena that non-expected utility models may be necessary to explain. Of course no theory is perfect, so might it be that these are exceptional cases and that EUT is still a reasonable approximation for a wide range of field behavior? I suspect this is

little more than wishful thinking. In a new paper, Matthew Rabin (forthcoming) presents a "calibration theorem" which shows that expected utility theory has some grossly implausible implications. The central result is that an expected utility maximizer who displays risk aversion in cases where outcomes are modest will display ludicrously high degrees of risk aversion over large stakes and the result holds for any concave utility function. Rabin argues that this general property of EUT is at odds with intuition and observation in both the lab and the field. For example, there is a huge amount of evidence which shows risk aversion over small stake laboratory choices, and explaining this using EUT implies pathologically risk averse behavior over larger outcomes. Rabin's argument should jolt us out of wishful thinking, since it suggests that EUT is implausible as a general account of behavior under risk.

5.3 *Theoretical Applications*

Since my concern is with theories as descriptive models, I have placed emphasis on assessing the predictive power of alternative theories. But while prediction is important, it is not everything. Other important questions surround the theoretical usefulness of alternatives to EUT. The standard theory is, without doubt, a potent simplification which can be easily applied in a range of theoretical contexts, and its use is pervasive. While a good deal of effort has been devoted to developing alternatives to EUT, by comparison, the use of such models in theoretical work outside of the specialist literature has been limited. Does this suggest that alternative models are too complex or intractable to be useful in a broader theoretical context? In general I think the answer is no and that other factors most likely explain the relatively slow take-up of new models. It is fair to say that giving up EUT raises some deep theoretical questions in fundamental areas of economics like game theory and the analysis of dynamic choices. For example, Nash equilibrium may fail to exist with non-EU preferences, and choices may be dynamically inconsistent. There are difficult problems to address here, but in mind of that perhaps we should expect progress to occur slowly and not take the gentle pace as evidence of intractability.

It is worth noting that many standard results and techniques are robust to some relaxations of the independence axiom: that is, economic theory as we know it does not simply implode when non-expected utility preferences are allowed. Although EUT has been a central building block in core areas of economics, many tools and results that have been developed assuming it actually require weaker assumptions. This was an important message of Machina's (1982) analysis which goes through even though his empirical hypothesis (that indifference curves fan out) was not supported by the data. His generalized expected utility analysis enables us to extend theoretical results and insights derived from EUT to a non-expected utility framework. For example, so long as preferences have the necessary smoothness properties, we can characterize risk aversion, stochastic dominance preference, and comparative risk aversion in terms of properties of local utility functions. Hence, much of our understanding of these aspects of risk preference remains intact for a wide class of non-expected utility models (see Machina 1987, 1989). Moreover, our understanding of risk aversion has been substantially refined by discussions of the concept in a non-expected utility framework. For instance, the rank dependent approach allows us to

decompose risk aversion into elements deriving from, respectively, attitudes to consequences and attitudes to chance (Wakker 1994); and as we have seen, such models also provide accounts of observed risk behavior inconsistent with received notions of risk preference.

In several areas of applied theory, papers have emerged showing that well-known results usually derived assuming EUT do not rely on independence. For example, in the theory of auctions betweenness is sufficient to guarantee value revealing behavior in ascending bid auctions (Karni and Safra 1989a); in a model of search discussed by Karni and Safra (1990), qualitatively similar stopping rules characterize the optimal behavior of EU maximizers and non-EU maximizers so long as the preferences of the latter are quasi-convex. Such correspondences, while interesting, are of course special cases. In general, the behavior of EU maximizers and non-EU maximizers does not coincide, so given the well-documented predictive failure of EUT, there is surely a good case for seeking to develop general tools of economic analysis capable of handling non-EU preferences. Such tools would then help us to understand what implications failures of EUT have for a wider class of economic phenomena. While that sounds like a major undertaking, some significant progress has already been made, and one example is in the area of game theory.

It is well-known that if players' preferences do not satisfy the independence axiom of EUT, Nash equilibrium may fail to exist. Independence is not necessary for existence: quasi-concavity or betweenness is enough. The problem case for standard game theory is quasi-convexity, and it is easy to understand the intuition behind this. In games where the only equilibria are mixed strategies, if players' preferences are quasi-convex, then they will be unwilling to randomize in the way required for equilibrium. Those interested only in normative analysis might brush this aside, but since we know that betweenness is not supported empirically, and violations of it often go in the direction of quasi-convexity (see Camerer and Ho 1994) a natural question for the descriptively minded economist is: how do we analyze strategic behavior when preferences are quasi-convex? Vincent Crawford (1990) provides an answer showing that a generalization of Nash, the "equilibrium in beliefs," which coincides with the standard concept for quasi-concave preferences, also exists when preferences are quasi-convex. Crawford's analysis makes an important theoretical step in showing how standard game theoretic tools can be extended to handle players with non-EU preferences. Other related discussions of non-expected utility theory in the context of games can be found in Karni and Safra (1989a,b) and Dekel, Safra, and Segal (1991).

It is true that giving up EUT has dramatic implications in some areas of theory, and one pertinent example is the area of dynamic choice. If EUT does not hold, then sequential choices may be *dynamically inconsistent*. To appreciate the significance of this, consider a sequential choice problem represented by a standard decision tree. An agent who is dynamically inconsistent may identify an optimal path viewed from the initial choice node, but then be unwilling to take actions that form part of that optimal path at choice nodes further down the tree. Wakker (1999) suggests an analogy between dynamic inconsistency and schizophrenia: the dynamically inconsistent agent has something akin to a split personality, with different aspects of the person revealing themselves in different parts of the tree. Although

some might regard this as a "problem" with non-expected utility models, I think that conclusion could be misleading for two reasons, one theoretical, the other empirical.

From the theoretical point of view it is important to note that relaxation of independence does not necessarily imply dynamic inconsistency. Machina (1989b) has shown that agents with non-expected utility preferences can be dynamically consistent if we are prepared to sacrifice the assumption of *consequentialism*. An implication of consequentialism in standard decision-tree analysis is that agents are entirely forward looking: at any given decision node, the consequentialist decision maker ignores any part of the tree that cannot be reached moving forward from that node. In contrast, Machina argues that risks born in the past may be relevant to current decisions and he provides some telling examples of where that could be the case. As such he defends the notion of a dynamically consistent non-EU agent by rejecting consequentialism.

It has only recently been properly understood that axioms of EUT, including the independence axiom, follow from assuming certain principles of dynamic choice (see Hammond 1988; Edward McClennen 1990; Robin Cubitt 1996). This provides a new form of normative defence for EUT. On the other hand, since we know that independence fails empirically, at least one of the dynamic choice principles that jointly imply it must be failing too. It follows that if we want to predict the behavior of real agents in dynamic contexts we will need models of dynamic decision making that relax the suspect dynamic choice principle(s) implicit in EUT. Several papers have investigated models of dynamic decision making that relax standard assumptions. Among the important contri-

butions are Machina (1989b), Karni and Safra (1989b, 1990), McClennen (1990) and Segal (1990, 1997). Since the models proposed by these authors give up different principles (for example, Machina and McClennen relax consequentialism; Segal (1990) relaxes the reduction of compound lotteries axiom) an obvious question to ask is: which principle or principles of dynamic choice are actually implicated when the independence axiom of EUT is violated? As yet, relatively little work has addressed this issue directly, though a recent experimental investigation by Cubitt, Starmer, and Sugden (1998a) suggests a surprising answer. In this experiment, common ratio type violations of independence appear to be due to the failure of a principle which we call timing independence. The answer is surprising since timing independence is implicit in most proposed models of dynamic decision making: the only exception I know of is Karni and Safra's (1989b, 1990) model of *behaviorally consistent choice*.

I bring this section to a close with a brief smorgasbord of applications. Non-expected utility models have been used across a reasonably diverse range of theoretical applications. Here are some examples using conventional approaches. Epstein and Zin (1989, 1991) use non-expected utility models in the context of intertemporal consumption and asset demand. Their approach allows the separation of parameters (i.e., for risk aversion, intertemporal substitution, and preference over the timing resolution of uncertainty) which are confounded in the conventional expected utility approach. Epstein (1995) discusses a range of applications of other non-expected utility models in macroeconomics, finance, and game theory. Epstein and Segal (1992) derive a social welfare function based on quadratic utility theory. Segal (1988b), Karni

(1995), and Machina (1995) consider the implications of non-expected utility for insurance. Chew (1985) and Neilson (1992) apply the model of implicit utility to, respectively, demand revelation in an auction context and asset demand.

The most popular non-expected utility models in applied work have been those based on the rank dependent form. Early applications include Yaari's (1987) analysis of portfolio selection; Chew, Karni, and Safra (1987) on risk aversion; Segal, Spivak, and Ziera (1988) on savings and risk aversion; Epstein and Zin (1990) on asset pricing in a recursive utility framework; plus Quiggin's (1991a,b; 1993) discussions of comparative statics, optimal lottery design, portfolio selection, and information revelation. More recently, the model has been applied to a diverse range of topics, including work by Chateauneuf and Cohen (1994) and Cohen (1995) on risk aversion; Eide (1995) on the effects of punishment in a model of crime; a series of papers investigating utility elicitation procedures (Wakker and Daniel Deneffe 1996; Han Bleichrodt and Quiggin 1997; and Bleichrodt, Jose Luis Pinto, and Wakker 1999); plus an analysis of dynamic decision making which allows a psychological role for anxiety in relation to unresolved lotteries (Wu 1999).

Attempts have also been made to apply nonconventional theories. The applications of prospect theory discussed in Section 5.2 (above), Thaler (1980) and Camerer (forthcoming), are relevant here. Also Jeffrey Rachlinski (1990, 1994, 1996) has applied prospect theory to the economics of litigation and civil negotiation. Applications of regret theory include Cubitt and Sugden's (1998) model of preference evolution (I discuss this below). Milton Weinstein and Robert Quinn (1983) propose a model of post-decisional "blame" which utilizes the notion of regret. A connected literature suggests that considerations of regret may be important particularly in the context of health-related decisions. For example, Weinstein (1986) considers whether regret *ought* to play a role in clinical decisions; René Richard (1994) presents survey evidence indicating that anticipation of post-decisional emotions, particularly regret, may be important factors in decisions relating to contraception. Richard argues that it may be possible to influence individual behavior through information campaigns that highlight potential regrets, and he discusses the implications for public policy in relation to, among other things, the spread of sexually transmitted diseases such as AIDS. Richard Smith (1996) proposes a modified version of regret to account for the valuation of health states. Not all of these applications involve formal economic models. They are none the less interesting. Indeed, there is a sense in which the formal models are lagging behind the less formal discussions. There is considerable evidence that considerations of regret do influence ordinary people in important decisions. At the same time, it seems clear that our existing models do not provide good formalization of such processes.

It has to be said that, overall, the volume of work applying non-expected utility models looks quite small given how long some of the theories have been available. I think things may be changing and that we will see increasing use of models based on the rank dependent form. Until recently, the sheer variety of competing models probably counted against their use. Too many alternatives were on offer with no obvious way to discriminate between them (bear in mind that many of these theories were proposed to explain the same, relatively small, set of choice anomalies).

But now that much more evidence has accumulated, it seems clear that there are quantitatively important phenomena that should not be ignored in general economic analysis. One of these is surely the phenomenon of nonlinear probability weighting. The rank dependent model is likely to become more widely used precisely because it captures this robust empirical phenomenon in a model which is quite amenable to application within the framework of conventional economic analysis.

Loss aversion is another empirically important concept, and I sense that economists are becoming more interested in studying the implications of assuming loss-averse preferences for a range of economic issues. Tversky and Kahneman (1991) present a model—based on prospect theory—that applies the ideas of reference dependence and loss aversion in *riskless* choice, and attempts are currently under way to examine the implications of rank dependent preferences for fundamental theoretical issues in economics. For example, Munro (1998) examines the implications for welfare economics of assuming reference dependent preferences; Munro and Sugden (1998) examine the conditions necessary for general equilibrium in an economy where agents have reference dependent preferences.

Sign and rank dependent models—like cumulative prospect theory—capture both of these empirically important phenomena in a theoretically compact way. And, while not all of the empirical evidence fits this approach, it does provide an account consistent with some of the most robust stylized facts from a range of experimental studies.[28] Since

these models are essentially conventional, and since their use seems to be expanding, general claims to the effect that they are intractable or not useful in economics more broadly seem unconvincing. Perhaps there is a case for thinking that the position we should now aim for is one in which models like cumulative prospect theory become the default in applied economics with EUT used as a convenient special case, but only when we can be confident that loss aversion and probability weighting are insignificant. While that position may be some way off, my prediction is that the use of models incorporating probability weights and loss aversion will grow rapidly, and my normative judgement is that, if it doesn't, it ought to.

6. *Phase III: New Directions in the Theory of Choice under Risk*

I have argued that there is an established case for taking non-expected utility models seriously and, further, taking unconventional approaches to modeling seriously. My arguments have, in several places, relied heavily upon experimental evidence to make the case. This strategy, however, begs an important question, which is: are the choices of badly behaved experimental subjects particularly, if at all, relevant to economic enquiry? The experimental paradigm is relatively new to economics and is still viewed with suspicion by some members of the discipline. Since most of the data driving developments in this area have come from experimental investigations, could it be that choice anomalies are revealing defects of the experimental method (in relation to economics) as opposed to flaws in conventional economic theory? I have argued elsewhere (Starmer 1999b) that it would be hard for economists to argue for a blunt rejection of experimental data per se. That said, there are grounds for meaningful

---

[28] For those interested in where rank-dependent models fail, aside from the cases discussed in Section 5.1.2 which count against all conventional theories, see Wakker, Erin, and Weber (1994), Wu (1994) and Birnbaum, Jamie Patton, and Melissa Lott (1999).

debate about the theoretical and empirical significance of laboratory-observed choice anomalies. Some economists have challenged the empirical significance of laboratory anomalies for economic investigation by suggesting that they may not generalize to economically meaningful contexts. Since laboratory experiments are usually designed to control precisely those variables that economic theories identify as important, the arguments along these lines that merit serious consideration, in my view, are those that offer *some* account of why behavior in laboratory contexts may not extend to contexts of general concern. I shall focus on two lines of argument of particular interest here as they are beginning to open up new and exciting theoretical accounts of individual choice behavior.

6.1 *The Evolution of Preference*

Most of the data we have been discussing has come from individual choice experiments where subjects undertake a series of one-off tasks. Typically both the tasks and the environment will be unfamiliar to subjects. Moreover, although most experiments involve real —usually monetary—incentives, the most common reward mechanism is the random lottery incentive system. In experiments with this design, subjects are rewarded according to their response to one task which is randomly selected at the end of the experiment.[29] This provides little or no opportunity for subjects to revise their behavior in the light of feedback on the consequences of prior choices. Some have argued that such "raw" behavior may have little in common with the behavior of agents in economic environments where there is opportunity

to learn. Of course, many important decisions are taken rarely and afford limited opportunity for repetition, change, learning, and so on. Nevertheless, it is interesting to ask whether the "anomalies" observed in laboratory behavior disappear in market contexts or where appropriate conditions for learning exist.

Charles Plott (1996) argues that, in the presence of sufficiently strong incentives, laboratory anomalies tend to disappear when subjects are allowed to adjust their behavior on the basis of experience gained through repetition. As Plott puts it: "Behavior seems to go through stages of rationality that begin with a type of myopia when faced with unfamiliar tasks. With incentives and practice, which take the form of repeated decisions in the experimental work (but might include play, banter, discussions with others, stages of commitment, etc.), the myopia gives way to what appears to be a stage of more considered choices that reflect stable attitudes or preference" (1996, p. 248). Plott calls this the *discovered preference hypothesis.* The argument is essentially empirical: he draws on a range of experimental evidence indicating a tendency for anomalous behavior to converge on the predictions of economic theory when choices are repeated in market-like settings. The collected evidence is, without doubt, impressive, but the vast majority of it relates to experiments in which preferences of the experimental subjects have been "induced" or controlled.[30] The purpose of such experiments is usually to investigate whether particular market forms generate the equilibria predicted by economic theory when we know what the subjects' preferences are. For our purposes, however, the more relevant question is whether agents who lack,

---

[29] For a discussion of this random lottery design, and evidence on its validity, see Starmer and Sugden (1991); Jane Beattie and Loomes (1997); and Cubitt, Starmer, and Sugden (1998b).

[30] For an account of the induced preference methodology see Smith (1976).

say, expected utility preferences might evolve behavior more consistent with them through market participation. There has been relatively little empirical investigation of how preferences themselves might be affected by exposure to market mechanisms. Plott discusses only one study with a direct bearing on this. He cites experimental work by James Cox and David Grether (1996) as support for the conclusion that "the classical preference reversal can be seen as a product of inexperience and lack of motivation, and it goes away with experience of a market setting" (Plott 1996, p. 231). Although there is some support for that claim in Cox and Grether, their results are complicated and do not provide the basis for a general assertion that anomalies like preference reversal tend to disappear in *any* economically relevant market context.

There is some evidence, albeit relatively limited, relating to whether other violations of expected utility persist in environments that might allow subjects' behavior to evolve. For example, Mikhail Myagkov and Plott (1997) find evidence consistent with prospect theory (e.g. implications of reference dependence and diminishing sensitivity) in an experimental market environment with repeated decisions. While the authors identify some tendency for behavior to change in a direction consistent with expected utility preferences, the data *do not* reveal general convergence on the predictions of the standard model. Even where violations of EUT do fall in market contexts, preference discovery may not be the right interpretation of the effect. For example, Dorla Evans (1997) finds that using a market mechanism to elicit valuations for gambles, rather than an individual pricing task, leads to a marked reduction in betweenness violations. But the effect is not due to any preferences having changed when the market is introduced; it occurs because the chosen market mechanism (a sealed-bid auction) happens to select a price in the middle of the distribution of bids. An experiment conducted by John Bone, John Hey, and John Suckling (1999) suggests that repetition and group discussion *increases* Allais-type violations of EUT. Thus, there is a good case for thinking that patterns of behavior change in some environments involving markets and/or repetition, but as yet there is no sound empirical basis for asserting a general tendency towards expected utility preferences under "market conditions." The evidence is at best mixed. Is there a theoretical basis for assuming that behavior may evolve towards expected utility (or anything else)?

For decisions that are rare and/or irreversible (e.g. childbearing, marriage, job-taking, decisions relating to health and education) the scope for learning must be limited. Repeated choices that generate feedback (e.g. stock market decisions, some consumption decisions) are the more obvious candidates. As yet economics lacks any well-developed theoretical account of learning, the impact of incentives, and so on. Moreover, at least some of the work that does exist shows there can be no presumption that learning will always generate behavior that converges on full-blown rationality (see for example Timermann 1994). There is, of course, a considerable and growing volume of research on evolutionary models in economics—a good part of this literature is reviewed by Richard Nelson (1995)—though so far there has been relatively little attempt to examine whether theories of preference may be based on evolutionary foundations. Exceptions are Karni and David Schmeidler (1986) who argue that the expected utility hypothesis may be derived from a principle of *self-preservation,* plus recent papers by Tilman

Börgers and Rajiv Sarin (1996, 1997) and Cubitt and Sugden (1998). These later contributions which examine models in which preferences evolve under the pressure of some selection mechanism reach more skeptical conclusions. Although the selection mechanisms considered by the two sets of authors are quite different—in Börgers and Sarin, the selection mechanism is reinforcement learning, in Cubitt and Sugden it is imitation—a common conclusion emerges: expected utility preferences evolve only under restrictive assumptions. Moreover, Cubitt and Sugden argue that imitation does not even imply pressure to select for transitive preferences over prospects. While there remains much to be done before we could claim to have an adequate understanding of the evolution of behavior in real market environments, already it seems clear that it would be rash to assume a general tendency for individual behavior in relation to risk to converge with conventional economic assumptions simply by virtue of it taking place within some real market institution.

6.2 *Theories of Stochastic Preference*

In the 1990s a new direction in the hunt for a descriptive theory has emerged with the publication of a number of papers investigating models of stochastic preference. This literature is generating new ways of modeling choice behavior and new methods for testing existing theories. It also provides a reason to re-examine some of the earlier conclusions drawn from experimental evidence.

The models we have discussed so far are essentially deterministic and, if interpreted literally each could be rejected by a single contrary observation. This seems too strong a test, and in most empirical work researchers have interpreted theories stochastically. The typical strategy has been to test whether observed violations of a theory can be explained as "random error" and theories are rejected only when the departure seems systematic (i.e., non-random). For example, in an experimental test for a common ratio effect in problems like those of Figure 2, an individual subject could violate expected utility in either of two ways: they could choose s_1^{**} then r_2^{**} (as predicted by fanning-out theories) or they could choose r_1^{**} then s_2^{**}. It would be common for researchers to adopt the null hypothesis that individuals choose according to expected utility plus random error and to further assume that random error implies the two violating patterns are equally likely. Hence, on this procedure, simply observing some people choosing consistently with fanning out is not enough to reject the null; the typical test examines whether violation consistent with the common ratio effect occurs significantly more frequently than its mirror image. Similarly, in testing for the cycle of choice predicted by regret theory, some researchers have operationalized a null hypothesis of EUT plus error by assuming that the probability of observing the predicted choice cycle equals that of observing a cycle in the opposite direction. While such assumptions lacked any substantive theoretical underpinnings, in the absence of any generally accepted theory of randomness, researchers had little choice but to base empirical tests on ad hoc assumptions about an error generating process.

A series of recent papers by Hey and Chris Orme (1994), Harless and Camerer (1994), and Loomes and Sugden (1995) has opened the way to a more general approach by suggesting different ways of modeling stochastic preferences. Each approach provides a general framework for developing stochastic versions of

alternative (deterministic) "core" theories of preference such as expected utility theory or some non-expected utility theory, though the interpretation of the source of randomness differs between the models. In the Hey and Orme model, the choice between a pair of prospects **q** and **r** is determined by the sign of :

$$HO = [V(\mathbf{q}) - V(\mathbf{r})] + \varepsilon$$

where $V(.)$ is the preference function of a deterministic core theory. If $\varepsilon = 0$, preferences are determined purely by the core theory, and if HO is positive (negative) then **q** (**r**) is chosen. Hey and Orme model the stochastic component by assuming that ε is a normal variate with a mean of zero. Consequently, the sign of HO can be reversed by the draw of a large enough ε with the "right" sign. This is essentially the *Fechner model* discussed earlier by Gordon Becker, Morris DeGroot, and Jacob Marchak (1963). Hey and Orme interpret the randomness in this model as some kind of calculation error. Notice that the larger the difference in values assigned to the prospects by the core theory, the less likely it is that true preferences will be overturned by the error term. Harless and Camerer's approach is different. They assume that any decision reveals true preferences (as defined by the core theory) with probability $1 - e$, but there is some constant probability e, that the individual chooses at random. The error generating mechanism here is akin to the *trembling hand* idea, familiar in game theory. In these first two models the stochastic element reflects deviations from "true" preferences resulting from miscalculations, slips or trembles, and so on. Loomes and Sugden consider a *random preference model*—also discussed by Becker, deGroot, and Marchak—which has a different interpretation: for any given choice, the individual acts on preferences satisfying the restrictions of the

core theory, but the parameters of the core theory to be applied to any given choice are determined by a random process. So, if the core theory is expected utility, the random draw determines the individual's degree of risk aversion, independently, for each choice. In this case, the stochastic element is inherent in preferences as opposed to random deviation about true preferences.

It is well-known that the data from choice experiments show a great deal of variability. For instance, a common finding is that individuals confronted with the same pairwise choice problem twice within a given experiment frequently give different responses on the two occasions.[31] Stochastic choice is more convincing than indifference as an account for such intrinsic variability, but explicitly introducing randomness into models of choice also provides alternative possible explanations for at least some of the violations of EUT that motivated non-expected utility models. For instance, assuming expected utility as the core theory, Loomes and Sugden show that all three error models allow the possibility of *systematic* cycles of choice. Thus the introduction of a stochastic component might allow the explanation of intransitive choice, without giving up the assumption of transitive preferences in the core theory. They also show that a model combining EUT with the HO specification can generate behavior consistent with fanning out (or fanning in) like the common ratio effect. Given this, it seems natural to ask just how much of the known data could be explained by a stochastic version of EUT.

An analysis of experimental choice

---

[31] A variety of studies including Starmer and Sugden (1989), Camerer (1989), Hey and Orme (1994), and T. Parker Ballinger and Nathaniel Wilcox (1997) find that between one-quarter and one-third of subjects "switch" preferences on repeated questions.

data conducted by Harless and Camerer (1994) sheds some light on this. They proposed their "tremble" theory as part of an econometric methodology designed to compare the predictive power of EUT and a variety of alternatives against experimental choice data while allowing for randomness. Their study, which examined data from tens of thousands of choices from over twenty prior studies, provides a highly comprehensive snap-shot of the evidence. The approach they take allows them to assess the relative performance of alternative theories given different trade-offs between parsimony (the number of patterns the theory allows) and predictive accuracy. With most emphasis placed on predictive accuracy (and least on parsimony) Harless and Camerer find that the "best" theory will be one that allows some mixed fanning in and fanning out; with the premium on parsimony, simple models like expected utility, or expected value, win. Interestingly, however, when choices involve prospects with "mixed support" (i.e. involve some choices on triangle boundaries), they find that there is no trade-off between parsimony and fit which picks expected utility as the best theory. In general they conclude that "The pairwise-choice studies suggest that violations of EU are robust enough that modeling of aggregate economic choice behavior based on alternatives to EU is well worth exploring" (p. 1287). In addition, they find that there is room for improvement in the predictive power of new theories: "For every theory there is systematic variation in excluded patterns which could in principle be explained by a more refined theory" (p. 1284).

Although this study is based on an impressive data base, it is important to recognize that their analysis involves specific assumptions about error, parsimony, and so on, to which the conclu-

sions may be sensitive. A parallel study by Hey and Orme (which appears back-to-back with Harless and Camerer in the 1994 volume of *Econometrica*) concludes with a more positive assessment of EUT. Their analysis, which uses the HO error specification, allows the best model to be determined separately for each individual (Harless and Camerer fit models to aggregate data). Hey and Orme find that EUT works as well as any other model for a substantial *minority* of their subjects (almost 40 percent). Even so, for the majority of their subjects at least one non-expected utility model outperforms the standard theory and the rank dependent model is a leading contender (along with the quadratic utility model).

Once we think in terms of alternative models of error, the problem of theory selection becomes more complex since the number of available models is now the product of the available core theories and error specifications. Harless and Camerer, and Hey and Orme, each compare models for a given error specification. A number of researchers have now begun to address the problem of choosing between error specifications. Loomes and Sugden (1998) and Enrica Carbone (1997) compare the predictive performance of EUT under each of the three error specifications. A tentative finding from this research seems to be that the trembling hand model performs relatively poorly. However, Loomes, Moffat, and Sugden (1998) investigate the possibility of two more sophisticated error theories that extend either the Fechner model or the random preference model to allow the possibility of trembles. They compare two core theories—expected utility and rank-dependent expected utility—and find that the data supports rank-dependent expected utility combined with a random preference error term *plus trembles*. They

report another finding too, which gives pause for thought. Their data derives from an experiment in which subjects were confronted with repeated choices. They find two significant trends in their data: (i) randomness due to trembles tended to decay with experience; and (ii) deviations from EUT decayed with experience. Is this evidence of individuals discovering expected utility preferences? I think the truth is that we don't yet know. There is growing evidence that there are some dynamic processes at work in repeated choice contexts. For example, Hey and Orme (1994), Ballinger and Wilcox (1997), and Loomes and Sugden (1998) all report systematic variation in risk aversion across repeated choices. In my view, an examination of those processes constitutes another extremely interesting avenue for future research.

7. *Concluding Thoughts*

In 1994 I attended a conference on the foundations of utility and risk[32] which concluded with a roundtable discussion where members of a panel, including Mark Machina and Robert Sugden, were asked to offer their personal reflections on the future of non-expected utility. Machina argued that while much good work had been done in developing theoretical alternatives to EUT, the field of applications was relatively underdeveloped. Perhaps the time had come, he suggested, to spend more time considering the implications of these alternatives for a wider collection of economic problems. Sugden, on the other hand, suggested that the gathering evidence shows actual choice behavior to be more complex than our models, and he voiced skepticism about the ability of conventional theorizing to provide an account of such complexity.

He thought a central theoretical problem remained to be solved before we could claim to have developed a reasonably general descriptive model of choice. In my view they were both right.

There is substantial evidence that EUT is likely to be descriptively misleading in at least some important contexts and, given the accumulating evidence supporting, in particular, probability weighting and loss aversion, we have at least some well-grounded hypotheses about important factors generating departures from the standard theory. From a practical point of view the rank-dependent models (including the sign-dependent variants) provide a convenient way of modeling these well established influences on choice, and there seems good reason to push forward the task of examining what implications such models have in general economic contexts.

At the same time, there seems to be a good case for pushing at the limits and perhaps stepping beyond the boundaries of conventional theorizing. In the past, some have justified the conventional approach on the grounds that its assumptions were supported by the evidence. It is hard to do that in any convincing way now, as it seems quite plain that real behavior refuses to be confined by the limits of conventional theorizing. We are discovering that a wide variety of behavior, in both the lab and the field, cannot be adequately explained within the conventional framework. While existing unconventional theories have their own limitations, models like regret and prospect theory (the original version) illustrate the possibility of working outside the conventional framework. Insights derived from these theories have proved useful in understanding real behavior even if the formal theories have not been widely applied. Moreover, the investigation of

[32] FUR VII, Oslo.

these theories has led to the discovery of new empirical phenomena, including, for example, the discovery of the cyclical choices predicted by prospect theory and the event-splitting effect. This must surely count as progress too.

Let me offer a final thought about the significance of the evidence that has been driving this research program. Perhaps "anomalies" in choice behavior will turn out to be insignificant for a broad range of economic problems. I have to confess that could be true. On the other hand, theories of choice lie at the very heart of economics, and the data we have suggests that choice behavior displays complex patterns in even very simple contexts. Given that datum line, is there a good reason for thinking that behavior in a more complex reality will better conform to our simple models? It would surely be genuinely surprising if that were true, and a major intellectual achievement if it could be demonstrated. Had it not been for the program of research investigating failures of EUT and alternative models of choice, we may not even have asked such questions. Before this work became well known, most economists probably took it for granted that EUT was the right model of individual choice. Even if it turns out that EUT can be supported using arguments based on learning or evolution, we will have learned something new and important, that is, *why and in what circumstances EUT applies.* Currently, one can point to some evidence that behavior becomes "better-behaved" by exposure to markets, but there is relatively little evidence for this and it is mixed. It seems to me that some of the most interesting current research relates to how behavior evolves over time and across different institutional settings. As yet we have only the vaguest appreciation of the mechanisms involved. My hope is that research on these topics will continue to be driven by an effort to confront our theories with the evidence, however challenging the data may be.

REFERENCES

Abdellaoui, Mohammed. 1998. "Parameter-Free Eliciting of Utilities and Probability Weighting Functions," working paper, GRID, ENS, Cachan, France.

———. 1999. "A Genuine Rank-Dependent Generalization of von Neumann-Morgenstern Expected Utility Theorem," working paper, GRID, ENS, Cachan, France.

Ali, Mukhtar M. 1977. "Probability and Utility Estimates for Racetrack Bettors," *J. Polit. Econ.* 85, pp. 803–15.

Allais, Maurice. 1953. "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine," *Econometrica*, 21, pp. 503–46.

———. 1979. "The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School," in *Expected Utility Hypotheses and the Allais Paradox.* M. Allais and O. Hagen, eds. Dordrecht: Reidel.

Anand, Paul. 1993. *Foundations of Rational Choice under Risk*. Oxford U. Press and Clarendon.

Ballinger, T. Parker and Nathaniel T. Wilcox. 1997. "Decisions, Error and Heterogeneity," *Econ. J.* 107:443, pp. 1090–105.

Bateman, Ian; Alistair Munro, Bruce Rhodes, Chris Starmer, and Robert Sugden. 1997. "A Test of the Theory of Reference Dependent Preferences," *Quart. J. Econ.* 112:2, pp. 479–505.

Battalio, Raymond C.; John. H. Kagel, and Komain Jiranyakul. 1990. "Testing between Alternative Models of Choice under Uncertainty: Some Initial Results," *J. Risk Uncertainty*, 3, pp. 25–50.

Beattie, Jane and Graham Loomes. 1997. "The Impact of Incentives upon Risky Choice Experiments," *J. Risk Uncertainty*, 14, pp. 149–62.

Becker, Gordon M.; Morris H. DeGroot, and Jacob Marschak. 1963. "Stochastic Models of Choice Behavior," *Behavioral Sci.* 8, pp. 41–55.

Becker, Joao L. and Rakesh K. Sarin. 1987. "Gamble Dependent Utility," *Manage. Sci.* 33, pp. 1367–82.

Bell, David. 1982. "Regret in Decision Making under Uncertainty," *Operations Res.* 20, pp. 961–81.

———. 1985. "Disappointment in Decision Making under Uncertainty," *Operations Res.* 33, pp. 1–27.

Benartzi, Shlomo and Richard H. Thaler. 1995. "Myopic Loss Aversion and the Equity Premium Puzzle," *Quart. J. Econ.* 110, pp. 73–92.

Bernoulli, Daniel. 1954. "Exposition of a New Theory on the Measurement of Risk," (original 1738) *Econometrica*, 22, pp. 23–36.

Birnbaum, Michael H. and Juan B. Navarrete. 1998. "Testing Descriptive Utility Theories: Violations of Stochastic Dominance and Cumulative Independence," *J. Risk Uncertainty,* 17, pp. 49–78.

Birnbaum, Michael H. and Laura A. Thompson. 1996. "Violations of Monotonicity in Choices between Gambles and Certain Cash," *Amer. J. Psych.* 109, pp. 501–23.

Birnbaum, Michael H.; Jamie N. Patton, and Melissa K. Lott. 1999. "Evidence against Rank-Dependent Utility Theories: Tests of Cumulative Independence, Interval Independence, Stochastic Dominance, and Transitivity," *Org. Behavior Human Decision Processes*, 77:1, pp. 44–83.

Bleichrodt, Han and John Quiggin. 1997. "Characterising QALYs under a General Rank-Dependent Utility Model," *J. Risk Uncertainty*, 15:2, pp. 151–65.

Bleichrodt, Han; Jose Luis Pinto, and Peter P. Wakker. 1999. "Using Descriptive Findings of Prospect Theory to Improve Prescriptive Applications of Expected Utility," plenary presentation at SPUDM, Mannheim.

Bone, John; John Hey, and John Suckling. 1999. "Are Groups More (or Less) Consistent than Individuals?" *J. Risk Uncertainty*, 18, pp. 63–81.

Borch, Karl H. 1974. *The Mathematical Theory of Insurance.* Lexington, MA: .Lexington Books.

Börgers, Tilman and Rajiv Sarin. 1996. "Naive Reinforcement Learning with Endogenous Aspirations," mimeo, U. College London.

———. 1997. "Learning through Reinforcement and Replicator Dynamics," *J. Econ. Theory,* 77, pp. 1–14.

Camerer, Colin F. 1989. "An Experimental Test of Several Generalised Utility Theories," *J. Risk Uncertainty*, 2:1, pp. 61–104.

———. 1992. "Recent Tests of Generalizations of Expected Utility Theory," in *Utility: Theories Measurement, and Applications.* Ward Edwards, ed. Norwell, MA: Kluwer, pp. 207–51.

———. 1995. "Individual Decision Making," in *Handbook of Experimental Economics.* J. Kagel and A. E. Roth, eds. Princeton: Princeton U. Press.

———. 1998. "Bounded Rationality in Individual Decision Making," *Experimental Econ.* 1, pp. 163–83.

———. Forthcoming. "Prospect Theory in the Wild: Evidence from the Field," in *Choices, Values and Frames.* Daniel Kahneman and Amos Tversky, eds. Cambridge: Russell Sage Foundation.

Camerer, Colin F. and Teck-Hua Ho. 1994. "Violations of the Betweenness Axiom and Nonlinearity in Probability," *J. Risk Uncertainty*, 8, pp. 167–96.

Carbone, Enrica. 1997. "Investigation of Stochastic Preference Theory Using Experimental Data," *Econ. Letters*, 57:3, pp. 305–11.

Chateauneuf, Alain and Michèle Cohen. 1994. "Risk Seeking with Diminishing Marginal Utility in a Non-Expected Utility Model," *J. Risk Uncertainty*, 9, pp. 77–91.

Chew, Soo Hong. 1983. "A Generalization of the Quasilinear Mean with Applications to the Measurement of Income Inequality and Decision Theory Resolving the Allais Paradox," *Econometrica*, 51, pp. 1065– 92.

———. 1985. "Implicit-Weighted Utility Theories, M-Estimators, and Non-Demand Revelation of Second-Price Auctions for an Uncertain Auctioned Object," working paper 15, Johns Hopkins U.

———. 1989. "Axiomatic Utility Theories with the Betweenness Property," *Annals Operational Res*. 19, pp. 273–98.

Chew, Soo Hong and Larry G. Epstein. 1989. "A Unifying Approach to Axiomatic Non-expected Utility Theories," *J. Econ. Theory*, 49, pp. 207–40.

Chew, Soo Hong; Larry G. Epstein, and Uzi Segal. 1991. "Mixture Symmetry and Quadratic Utility," *Econometrica,* 59, pp. 139–63.

Chew, Soo Hong, Larry G. Epstein, and Peter P. Wakker. 1993. "A Unifying Approach to Axiomatic Non-Expected Utility Theories: Correction and Comment," *J. Econ. Theory*, 59, pp. 183–88.

Chew, Soo Hong; Edi Karni, and Zvi Safra. 1987. "Risk Aversion in the Theory of Expected Utility with Rank-dependent Probabilities," *J. Econ. Theory*, 42, pp. 370–81.

Chew, Soo Hong and Kenneth MacCrimmon. 1979a. "Alpha-nu Choice Theory: a Generalisation of Expected Utility Theory," working paper 669, U. British Columbia.

———. 1979b. "Alpha Utility Theory, Lottery Composition, and the Allais Paradox," Faculty of Commerce and Business Admin. working paper 686, U. British Columbia.

Chew, Soo Hong and William S. Waller. 1986. "Empirical Tests of Weighted Utility Theory," *J. Math. Psych.* 30, pp. 55–72.

Cohen, Michèle. 1995. "Risk-aversion Concepts in Expected- and Non-expected Utility Models," *Geneva Papers Risk Insurance Theory*, 20, pp. 73–91.

Conlisk, John. 1989. "Three Variants on the Allais Example," *Amer. Econ. Rev.* 79, pp. 392–407.

———. 1996. "Why Bounded Rationality," *J. Econ. Lit.* 34:2, pp. 669–700.

Cox, James C. and David M. Grether. 1996. "The Preference Reversal Phenomenon: Response Mode, Markets and Incentives," *Econ. Theory*, 7, pp. 381–405.

Crawford, Vincent, P. 1990. "Equilibrium without Independence," *J. Econ. Theory*, 50, pp. 127–54.

Cubitt, Robin P. 1996. "Rational Dynamic Choice and Expected Utility Theory," *Oxford Econ. Papers*, 48, pp. 1–19.

Cubitt, Robin P.; Chris Starmer, and Robert Sugden. 1998a. "Dynamic Choice and the Common Ratio Effect: An Experimental Investigation," *Econ. J.* 108:450, pp. 1362–80.

———. 1998b. "On the Validity of the Random Lottery Incentive System," *Experimental Econ.* 1, pp. 115–31.

Cubitt, Robin and Robert Sugden. 1998. "The Selection of Preferences through Imitation," *Rev. Econ. Stud.* 65, pp. 761–71.

Dekel, Eddie. 1986. "An Axiomatic Characterization of Preferences under Uncertainty: Weakening the Independence Axiom," *J. Econ. Theory*, 40, pp. 304–18.

Dekel, Eddie.; Zvi Safra, and Uzi Segal. 1991. "Existence and Dynamic Consistency of Nash Equilibrium with Non-expected Utility Preferences," *J. Econ. Theory,* 55, pp. 229–46.

Diecidue, Enrico and Peter P. Wakker. 1999. "On the Intuition of Rank-dependent Expected Utility," paper presented at FUR IX, Marrakesh.

Edwards, Ward. 1955. "The Prediction of Decisions among Bets," *J. Experimental Psych.* 50, pp. 201–14.

———. 1962. "Subjective Probabilities Inferred from Decisions," *Psych. Rev.* 69, pp. 109–35.

Eide, Erling 1995. "RDEU Models of Crime," Law and Economics working paper C-1, Universitetet i Oslo.

Epstein, Larry G. 1992. "Behavior under Risk: Recent Developments in Theory and Applications," in *Advances in Economic Theory: Sixth World Congress, Vol. 2.* Jean-Jacques Laffont, ed. Cambridge: Cambridge U. Press.

Epstein, Larry G. and Uzi Segal. 1992. "Quadratic Social Welfare Functions," *J. Polit. Econ.* 100, pp. 691–712.

Epstein, Larry G. and Stanley E. Zin. 1989. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, pp. 937–69.

———. 1990. "'First-order' Risk Aversion and the Equity Premium Puzzle," *J. Monetary Econ.* 26, pp. 387–407.

———. 1991. "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *J. Polit. Econ.* 99, pp. 263–86.

Evans, Dorla A. 1997. "The Role of Markets in Reducing Expected Utility Violations," *J. Polit. Econ.* 105, pp. 622–36.

Fishburn, Peter C. 1978. "On Handa's 'New Theory of Cardinal Utility' and the Maximization of Expected Return," *J. Polit. Econ.* 86, pp. 321–24.

———. 1982. "Nontransitive Measurable Utility," *J. Math. Psych.* 26, pp. 31–67.

———. 1983. "Transitive Measurable Utility," *J. Econ. Theory,* 31, pp. 293–317.

———. 1988. *Nonlinear Preference and Utility Theory.* Brighton: Wheatsheaf Books.

Gigliotti, Garry and Barry Sopher. 1993. "A Test of Generalized Expected Utility Theory," *Theory and Decision*, 35, pp. 75–106.

Gonzalez, Richard and George Wu. 1999. "On the Shape of the Probability Weighting Function," *Cognitive Psych.* 38, pp. 129–66.

Green, Jerry and Bruno Jullien. 1988. "Ordinal Independence in Nonlinear Utility Theory," *J. Risk Uncertainty*, 1, pp. 355–87.

Gul, Faruk. 1991. "A Theory of Disappointment in Decision Making under Uncertainty," *Econometrica*, 59, pp. 667–86.

Hagen, Ole. 1979. "Towards a Positive Theory of Preferences under Risk" in *Expected Utility Hypotheses and the Allais Paradox*. M. Allais and O. Hagen, eds. Dordrecht: Reidel.

Handa, Jagdish. 1977. "Risk, Probability, and a New Theory of Cardinal Utility," *J. Polit. Econ.* 85, pp. 97–122.

Hammond, Peter. 1988. "Consequentialist Foundations for Expected Utility," *Theory and Decision*, 25, pp. 25–78.

Harless, David W. 1992. "Predictions about Indifference Curves inside the Unit Triangle: a Test of Variants of Expected Utility," *J. Econ. Behavior Org.* 18, pp. 391–414.

Harless, David W. and Colin F. Camerer. 1994. "The Predictive Utility of Generalized Expected Utility Theories," *Econometrica*, 62, pp. 1251–89.

Hausman, Daniel. 1992. *The Inexact and Separate Science of Economics*. Cambridge U. Press.

Heath, Chip; Steven Huddart, and Mark Lang. 1999. "Psychological Factors and Stock Option Exercise," *Quart. J. Econ.* 114:2, pp. 601–27.

Herne, Kaisa. 1998. "Testing the Reference-Dependent Model: An Experiment on Asymmetrically Dominated Reference Points," *J. Behavioral Decision Making,* 11, pp. 181–92.

Hey, John D. and Chris Orme. 1994. "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62, pp. 1291–326.

Holt, Charles A. 1986. "Preference Reversals and the Independence Axiom," *Amer. Econ. Rev.* 76, pp. 508–15.

Humphrey, Stephen J. 1995. "Regret Aversion or Event-Splitting Effects? More Evidence under Risk and Uncertainty," *J. Risk Uncertainty*, 11, pp. 263–74.

———. 1998. "More Mixed Results on Boundary Effects," *Econ. Letters*, 61, pp. 79–84.

Jullien, Bruno and Bernard Salanié. 1997. "Estimating Preferences under Risk: The Case of Racetrack Bettors," IDEI and GREMAQ Toulouse U. working paper.

Kahneman, Daniel and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision under Risk," *Econometrica,* 47:2, pp. 263–91.

Karmarkar, Uday S. 1978. "Subjectively Weighted Utility: A Descriptive Extension of the Expected Utility Model," *Org. Behavior Human Performance,* 21, pp. 61–72.

Karni, Edi. 1995. "Non-expected Utility and the Robustness of the Classical Insurance Paradigm: Discussion," *Geneva Papers Risk Insurance Theory*, 20, pp. 51–56.

Karni, Edi and Zvi Safra. 1987. "'Preference Reversal' and the Observability of Preferences by Experimental Methods," *Econometrica*, 55, pp. 675–85.

———. 1989a. "Dynamic Consistency, Revelations in Auctions and the Structure of Preferences," *Rev. Econ. Stud.* 56, pp. 421–34.

———. 1989b. "Ascending Bid Auctions with Behaviorally Consistent Bidders," *Annals Operational Res.* pp. 435–46.

———. 1990. "Behaviorally Consistent Optimal Stopping Rules," *J. Econ. Theory*, 51, pp. 391–402.

Karni, Edi and David Schmeidler. 1986. "Self-Preservation as a Foundation of Rational Behavior Under Risk," *J. Econ. Behavior Org.* 7, pp. 71–81.

Knetsch, Jack L. 1989. "The Endowment Effect and Evidence of Non-reversible Indifference Curves," *Amer. Econ. Rev.* 79, pp. 1277–84.

Knetsch, Jack L. and J. A. Sinden. 1984. "Willingness to Pay and Compensation Demanded: Experimental Evidence of an Unexpected Disparity in Measures of Value," *Quart. J. Econ.* 99, pp. 507–21.

Lattimore, Pamela K.; Joanna R. Baker, and Ann D. Witte. 1992. "The Influence of Probability on Risky Choice: A Parametric Examination," *J. Econ. Behavior Org.* 17, pp. 377–400.

Lavoie, Marc. 1992. *Foundations of Post-Keynesian Economic Analysis*. Aldershot, UK: Edward Elgar.

Leland, J.W. 1998. "Similarity Judgements in Choice under Uncertainty: A Reinterpretation of the Predictions of Regret Theory," *Manage. Sci.* 44, pp. 659–72.

Lichtenstein, Sarah and Paul Slovic. 1971. "Reversals of Preference between Bids and Choices in Gambling Decisions," *J. Experimental Psych.* 89, pp. 46–55.

Lindman, Harold R. 1971. "Inconsistent Preferences among Gambles," *J. Experimental Psych.* 89, pp. 390–97.

Loewenstein, George F. and Daniel Adler. 1995. "A Bias in the Prediction of Tastes," *Econ. J.* 105, pp. 929–37.

Loomes, Graham; Peter Moffat, and Robert Sugden. 1998. "A Microeconometric Test of Alternative Stochastic Theories of Risk Choice," ERC discussion paper 9806, U. East Anglia.

Loomes, Graham; Chris Starmer; and Robert Sugden. 1989. "Preference Reversal: Information-Processing Effect or Rational Non-transitive Choice?" *Econ. J.* 99:395, pp. 140–51.

———. 1991. "Observing Violations of Transitivity by Experimental Methods," *Econometrica*, 59:2, pp. 425–39.

———. 1992. "Are Preferences Monotonic? Testing Some Predictions of Regret Theory," *Economica*, 59, pp. 17–33.

Loomes, Graham and Robert Sugden. 1982. "Regret Theory: An Alternative Theory of Rational Choice under Uncertainty," *Econ. J.* 92, pp. 805–24.

———. 1983. "A Rationale for Preference Reversal," *Amer. Econ. Rev.* 73, pp. 428–32.

———. 1986. "Disappointment and Dynamic Consistency in Choice under Uncertainty," *Rev. Econ. Stud.* 53:2, pp. 271–82.

———. 1987. "Some Implications of a More General Form of Regret Theory," *J. Econ. Theory*, 41:2, pp. 270–87.

———. 1995. "Incorporating a Stochastic Element into Decision Theories," *European Econ. Rev.* 39, pp. 641–48.

———. 1998. "Testing Different Stochastic Specifications of Risky Choice," *Economica*, 65, pp. 581–98.

Loomes, Graham and Caron Taylor. 1992. "Non-transitive Preferences over Gains and Losses," *Econ. J.* 102, pp. 357–65.

Luce, R. Duncan and Peter C. Fishburn. 1991. "Rank and Sign-dependent Linear Utility Models for Finite First-Order Gambles" *J. Risk Uncertainty*, 4, pp. 29–59.

MacCrimmon, Kenneth and Stig Larsson. 1979. "Utility Theory: Axioms versus Paradoxes," in *Expected Utility Hypotheses and the Allais Paradox*. M. Allais and O. Hagen, eds. Dordrecht: Reidel

Machina, Mark J. 1982. "'Expected Utility' Theory without the Independence Axiom," *Econometrica, 50, pp. 277–323.*

———. 1983. "The Economic Theory of Individual Behavior toward Risk: Theory, Evidence, and New Directions," Technical Report 433, Dept. Economics, Stanford U.

———. 1987. "Choice under Uncertainty: Problems Solved and Unsolved," *J. Econ. Perspectives*, 1, pp. 121–54.

———. 1989a. "Comparative Statics and Non-Expected Utility Preferences," *J. Econ. Theory*, 47, pp. 393–405.

———. 1989b. "Dynamic Consistency and Non-Expected Utility Models of Choice under Uncertainty," *J. Econ. Lit.* 27, pp. 1622–68.

———. 1994. "Review of 'Generalized Expected Utility Theory: The Rank-Dependent Model'" *J. Econ. Lit.* 32:3, pp. 1237–38.

———. 1995. "Non-Expected Utility and the Robustness of the Classical Insurance Paradigm," *Geneva Papers Risk Insurance Theory*, 20, pp. 9–50.

Marschak, Jacob. 1950. "Rational Behavior, Uncertain Prospects, and Measurable Utility," *Econometrica*, 18, pp. 111–41.

McClennen, Edward F. 1990. *Rationality and Dynamic Choice: Foundational Explorations.* Cambridge: Cambridge U. Press.

Mehra, Rajnish and Edward C. Prescott. 1985. "The Equity Premium Puzzle," *J. Monetary Econ.* 15, pp. 145–61.

McNeil, B. J.; S. G. Pauker, H. C. Sox, and Amos Tversky. 1982. "On the Elicitation of Preferences for Alternative Therapies," *New England J. Medicine*, 306, pp. 1259–62.

Moskowitz, Herbert. 1974. "Effects of Problem Presentation and Feedback on Rational Behavior in

Allais and Morlat-Type Problems," *Decision Sciences*, 5, pp. 225–42.

Mossin, J. 1968. "Aspects of Rational Insurance Purchasing," *J. Polit. Econ.* 76, pp. 553–68.

Munro, Alistair. 1998. "Loss Aversion and Welfare Economics," Economics Research Centre working paper, U. East Anglia.

Munro, Alistair and Robert Sugden. 1998. "A Theory of General Equilibrium with Reference Dependent Preferences," mimeo, U. East Anglia.

Myagkov, Mikhail and Charles R. Plot. 1997. "Exchange Economies and Loss Exposure: Experiments Exploring Prospect Theory and Competitive Equilibria in Market Environments," *Amer. Econ. Rev.* 87, pp. 801–28.

Neilson, William S. 1992. "A Mixed Fan Hypothesis and its Implications for Behavior toward Risk," *J. Econ. Behavior Org.* 19, pp. 197–211.

———. 1979. "Some Mixed Results on Boundary Effects," *Econ. Letters*, 39, pp. 275–78.

Nelson, Richard R. 1995. "Recent Evolutionary Theorizing about Economic Change," *J. Econ. Lit.* 33:1, pp. 48–90.

Neumann, John von and Oskar Morgenstern. 1947. *The Theory of Games and Economic Behaviour*. 2nd ed. Princeton: Princeton U. Press.

Payne, John W.; James R. Bettman, and Eric J. Johnson. 1993. *The Adaptive Decision Maker*. Cambridge U. Press.

Pidgeon, Nick. C. Hood, D. Jones, B. Turner and R. Gibson. 1992. "Risk Perception," in *Risk: Analysis, Perceptions and Management*. Report of a Royal Society Study Group. London: The Royal Society.

Plott, Charles, R. 1996. "Rational Individual Behavior in Markets and Social Choice Processes: The Discovered Preference Hypothesis," in *Rational Foundations of Economic Behavior*. K. Arrow, E. Colombatto, M. Perleman, and C. Schmidt, eds. London: Macmillan and NY: St. Martin's, pp. 225–50.

Prelec, Drazen. 1998. "The Probability Weighting Function," *Econometrica,* 66, pp. 497–527.

Preston, Malcolm G. and Phillip Baratta. 1948. "An Experimental Study of the Auction-Value of an Uncertain Outcome," *Amer. J. Psych.* 61, pp. 183–93.

Quiggin, John. 1982. "A Theory of Anticipated Utility," *J. Econ. Behavior Org.* 3:4, pp. 323–43.

———. 1991a. "Comparative Statics for Rank-Dependent Expected Utility Theory," *J. Risk Uncertainty*, 4, pp. 339–50.

———. 1991b. "On the Optimal Design of Lotteries," *Economica*, 58:229, pp. 1–16

———. 1993. *Generalized Expected Utility Theory*. Dordrecht: Kluwer.

———. 1994. "Regret Theory with General Choice Sets," *J. Risk Uncertainty*, 8:2, pp. 153–65.

Rabin, Matthew. Forthcoming. "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica*.

Rachlinski, Jeffrey J. 1990. "Prospect Theory and Civil Negotiation," working paper 17, Center on Conflict and Negotiation, Stanford U.

———. 1994. "Prospect Theory and the Economics of Litigation," PhD thesis, Stanford U.

———. 1996. "Gains, Losses, and the Psychology of Litigation," *South. Cal. Law Rev.*, 70, pp. 113–85.

Richard, René. 1994. "Regret Is What You Get: The Impact of Anticipated Feelings and Emotions on Human Behaviour," PhD thesis, Psych. Faculty, U. Amsterdam.

Rubinstein, Ariel. 1988. "Similarity and Decision-Making Under Risk," *J. Econ. Theory*, 46, pp. 145–53.

Samuelson, William F. and Richard Zeckhauser. 1988. "Status Quo Bias in Decision Making," *J. Risk Uncertainty*, 1, pp. 7–59.

Savage, Leonard J. 1954. *The Foundations of Statistics*. NY: Wiley.

Schoemaker, Paul. 1982. "The Expected Utility Model: Its Variants, Purposes, Evidence and Limitations," *J. Econ. Lit.* 20:2, pp. 529–63.

Segal, Uzi. 1988a. "Does the Preference Reversal Phenomenon Necessarily Contradict the Independence Axiom?" *Amer. Econ. Rev.* 78, pp. 233–36.

———. 1988b. "Probabilistic Insurance and Anticipated Utility," *J. Risk Insurance*, 55, pp. 287–97.

———. 1989. "Anticipated Utility: A Measure Representation Approach," *Annals Operational Res.* 19, pp. 359–73.

———. 1990. "Two-Stage Lotteries without the Independence Axiom," *Econometrica*, 58, pp. 349–77.

———. 1993. "The Measure Representation: A Correction," *J. Risk Uncertainty*, 6, pp. 99–107.

———. 1997. "Dynamic Consistency and Reference Points," *J. Econ. Theory*, 72, pp. 208–19.

Segal, Uzi and Avia Spivak. 1990. "First Order versus Second Order Risk Aversion," *J. Econ. Theory*, 51, pp. 111–125.

Segal, Uzi; Avia Spivak, and J. Ziera. 1988. "Precautionary Savings and Risk Aversion: An Anticipated Utility Approach," *Econ. Letters*, 27, pp. 223–27.

Simon, Herbert, A. 1955. "A Behavioral Model of Rational Choice," *Quart. J. Econ.* 69, pp. 99–118.

Slovic, Paul. 1995. "The Construction of Preferences," *Amer. Psychologist*, 50, pp. 364–71.

Slovic, Paul and Amos Tversky. 1974. "Who Accepts Savage's Axiom?" *Behavioral Sci.* 19, pp. 368–73.

Smith, Richard D. 1996. "Is Regret Theory an Alternative Basis for Estimating the Value of Health Care Interventions?" *Health Policy*, 37, pp. 105–15.

Smith, Vernon. 1976. "Experimental Economics: Induced Value Theory," *Amer. Econ. Rev. Proceedings*, pp. 274–79.

Starmer, Chris. 1992. "Testing New Theories of Choice under Uncertainty Using the Common Consequence Effect," *Rev. Econ. Stud.* 59, pp. 813–30.

———. 1999a. "Cycling with Rules of Thumb: An

Experimental Test for a New Form of Non-Transitive Behaviour," *Theory and Decision*, 46, pp. 141–58.

———. 1999b. "Experiments in Economics: Should We Trust the Dismal Scientists in White Coats?," *J. Econ. Method.* 6, pp. 1–30.

Starmer, Chris and Robert Sugden. 1989. "Violations of the Independence Axiom in Common Ratio Problems: An Experimental Test of Some Competing Hypotheses," *Annals Operational Res.* 19, pp. 79–102.

———. 1991. "Does the Random-Lottery Incentive System Elicit True Preferences?" *Amer. Econ. Rev.* 81, pp. 971–79.

———. 1993. "Testing for Juxtaposition and Event Splitting Effects," *J. Risk Uncertainty*, 6, pp. 235–54.

———. 1998. "Testing Alternative Explanations of Cyclical Choices," *Economica*, 65:259, pp. 347–61.

Stevenson, M. K.; J. R. Busmeyer, and J. C. Naylor. 1991. "Judgement and Decision-Making Theory," in *New Handbook of Industrial Organizational Psychology*. M. Dunnette and L. N. Hough, eds. Palo Alto, CA: Consulting Psychology Press.

Sugden, Robert. 1986. "New Developments in the Theory of Choice Under Uncertainty," *Bul. Econ. Res.* 38, pp. 1–24.

———1991. "Rational Choice: A Survey of Contributions from Economics and Philosophy," *Econ. J.* 101, pp. 751–85.

———. 1993. "An Axiomatic Foundation for Regret Theory," *J. Econ. Theory*, 60, pp. 159–80.

Tammi, Timo. 1997. "Essays on the Rationality of Experimentation in Economics: The Case of Preference Reversal," PhD thesis, U. Joensuu, Finland.

Thaler, Richard H. 1980. "Toward a Positive Theory of Consumer Choice," *J. Econ. Behavior Org.* 1, pp. 39–60.

Thaler, Richard H. and William T. Ziemba. 1988. "Parimutuel Betting Markets: Racetracks and Lotteries," *J. Econ. Perspectives*, 2, pp. 161–74.

Timmerman, Allan. 1994. "Can Agents Learn to Form Rational Expectations? Some Results on Convergence and Stability of Learning in the UK Stock Market," *Econ. J.* 104, pp. 777–97.

Tversky, Amos. 1969. "Intransitivity of Preferences," *Psych. Rev.* 76, pp. 31–48.

Tversky, Amos and Daniel. Kahneman. 1981. "The Framing of Decisions and the Psychology of Choice," *Science*, 211, pp. 453–58.

———. 1986. "Rational Choice and the Framing of Decisions," *J. Bus.* 59:4, pt. 2, pp. s251–78.

———. 1991. "Loss Aversion in Riskless Choice: A Reference-Dependent Model," *Quart. J. Econ.* 106:4, pp. 1039–61.

———. 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *J. Risk Uncertainty*, 5:4, pp. 297–323.

Tversky, Amos and Derek J. Koehler. 1994. "Support Theory: A Nonextensional Representation of Subjective Probability," *Psych. Rev.* 101, pp. 547–67.

Tversky, Amos; Paul Slovic, and Daniel Kahneman. 1990. "The Causes of Preference Reversal," *Amer. Econ. Rev.* 80:1, pp. 204–17.

Tversky, Amos and Richard H. Thaler. 1990. "Preference Reversals," *J. Econ. Perspectives*, 4, pp. 201–11.

Tversky, Amos and Peter P. Wakker. 1995. "Risk Attitudes and Decision Weights," *Econometrica,* 63, pp. 1255–80.

Viscusi, W. Kip. 1989. "Prospective Reference Theory: Toward an Explanation of the Paradoxes," *J. Risk Uncertainty,* 2, pp. 235–64.

Wakker, Peter P. 1994. "Separating Marginal Utility and Probabilistic Risk Aversion," *Theory and Decision*, 36, pp. 1–44.

———. 1999. "Justifying Bayesianism by Dynamic Decision Principles," plenary paper presented at FUR IX, Marrakesh.

Wakker, Peter P. and Daniel Deneffe. 1996. "Eliciting von Neumann–Morgenstern Utilities When Probabilities Are Distorted or Unknown," *Manage. Sci.* 42, pp. 1131–50.

Wakker, Peter P.; Ido Erev, and Elke U. Weber. 1994. "Comonotonic Independence: The Critical Test between Classical and Rank-Dependent Utility Theories," *J. Risk Uncertainty*, 9, pp. 195–230.

Wakker, Peter P.; Richard H. Thaler, and Amos Tversky. 1997. "Probabilistic Insurance," *J. Risk Uncertainty*, 15, pp. 7–28.

Weber, Martin; Franz Eisenführ, and Detlof von Winterfeldt. 1988. "The Effects of Splitting Attributes on Weights in Multi Attribute Utility Measurement," *Manage. Sci.* 34, pp. 431–45.

Weinstein, Milton, C. 1986. "Risky Choices in Medical Decision Making: A Survey," *Geneva Papers Risk Insurance*, 11, pp. 197–216.

Weinstein, Milton C. and Robert J. Quinn. 1983. "Psychological Considerations in Valuing Health Risk Reductions," *Nat. Resources J.* 23, pp. 659–73.

Wu, George. 1994. "An Empirical Test of Ordinal Independence," *J. Risk Uncertainty,* 9, pp. 39–60.

———. 1999. "Anxiety and Decision Making with Delayed Resolution of Uncertainty," *Theory and Decision*, 46, pp. 159–98.

Wu, George and Richard Gonzalez. 1996. "Curvature of the Probability Weighting Function," *Management Sci.* 42, pp. 1676–90.

Yaari, Menahem E. 1987. "The Dual Theory of Choice Under Risk," *Econometrica*, 55, pp. 95–115.