Choice Under Uncertainty: Problems Solved and Unsolved

Mark J. Machina

Choice Under Uncertainty:
Problems Solved and Unsolved

Mark J. Machina

F ifteen years ago, the theory of choice under uncertainty could be considered one of the "success stories" of economic analysis: it rested on solid axiomatic foundations, it had seen important breakthroughs in the analytics of risk, risk aversion and their applications to economic issues, and it stood ready to provide the theoretical underpinnings for the newly emerging "information revolution" in economics.[1] Today choice under uncertainty is a field in flux: the standard theory is being challenged on several grounds from both within and outside economics. The nature of these challenges, and of our profession's responses to them, is the topic of this paper.

The following section provides a brief description of the economist's canonical model of choice under uncertainty, the expected utility model of preferences over random prospects. I shall present this model from two different perspectives. The first perspective is the most familiar, and has traditionally been the most useful for addressing standard economic questions. However the second, more modern perspective will be the most useful for illustrating some of the problems which have beset the model, as well as some of the proposed responses.

Each of the subsequent sections is devoted to one of these problems. All are important, some are more completely "solved" than others. In each case I shall begin with an example or description of the phenomenon in question. I shall then review the empirical evidence regarding the uniformity and extent of the phenomenon. Finally, I shall report on how these findings have changed, or are likely to change, or ought to

---

■ *Mark J. Machina is Associate Professor of Economics, University of California, San Diego, La Jolla, California.*

change, the way we view and model economic behavior under uncertainty. On this last topic, the disclaimer that "my opinions are my own" has more than the usual significance.[2]

The Expected Utility Model

The Classical Perspective: Cardinal Utility and Attitudes Toward Risk

In light of current trends toward generalizing this model, it is useful to note that the expected utility hypothesis was itself first proposed as an alternative to an earlier, more restrictive theory of risk-bearing. During the development of modern probability theory in the 17th century, mathematicians such as Blaise Pascal and Pierre de Fermat assumed that the attractiveness of a gamble offering the payoffs $(x_1, \ldots, x_n)$ with probabilities $(p_1, \ldots, p_n)$ was given by its expected value $\bar{x} = \Sigma x_i p_i$. The fact that individuals consider more than just expected value, however, was dramatically illustrated by an example posed by Nicholas Bernoulli in 1728 and now known as the *St. Petersburg Paradox*:

> Suppose someone offers to toss a fair coin repeatedly until it comes up heads, and to pay you $1 if this happens on the first toss, $2 if it takes two tosses to land a head, $4 if it takes three tosses, $8 if it takes four tosses, etc. What is the largest sure gain you would be willing to forgo in order to undertake a single play of this game?

Since this gamble offers a 1/2 chance of winning $1, a 1/4 chance of winning $2, etc., its expected value is $(1/2) \cdot \$1 + (1/4) \cdot \$2 + (1/8) \cdot \$4 + \cdots = \$1/2 + \$1/2 + \$1/2 + \cdots = \$\infty$, so it should be preferred to any finite sure gain. However, it is clear that few individuals would forgo more than a moderate amount for a one-shot play. Although the unlimited financial backing needed to actually make this offer is somewhat unrealistic, it is not essential for making the point: agreeing to limit the game to at most one million tosses will still lead to a striking discrepancy between most individuals' valuations of the modified gamble and its expected value of $500,000.

The resolution of this paradox was proposed independently by Gabriel Cramer and Nicholas's cousin Daniel Bernoulli (Bernoulli, 1738/1954). Arguing that a gain of $200 was not necessarily "worth" twice as much as a gain of $100, they hypothesized that the individual possesses what is now termed a *von Neumann-Morgenstern utility function* $U(\cdot)$, and rather than using expected value $\bar{x} = \Sigma x_i p_i$, will evaluate gambles on the basis of expected utility $\bar{u} = \Sigma U(x_i)p_i$. Thus the sure gain ξ which would yield the same utility as the Petersburg gamble, i.e. the certainty equivalent of this gamble,

---

[2]In keeping with the spirit of this journal, references have been limited to the most significant examples of and/or most useful introductions to the literature in each area. For further discussions of these issues see Arrow (1982), Machina (1983a, 1983b), Sugden (1986) and Tversky and Kahneman (1986).
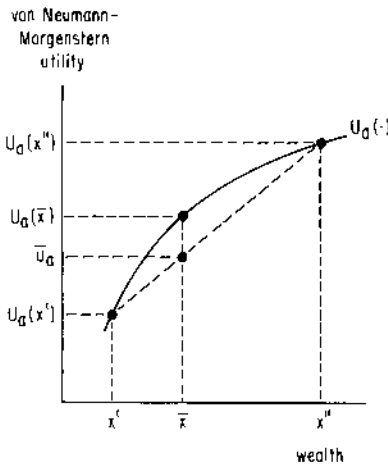
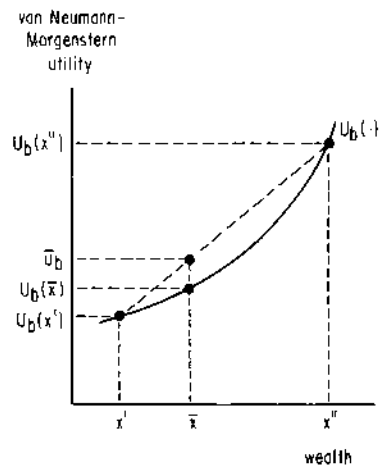Fig. 1a. Concave utility function of a risk averter

Fig. 1b. Convex utility function of a risk lover

is determined by the equation

$$(1) \qquad U(W + \xi) = (1/2) \cdot U(W + 1) + (1/4) \cdot U(W + 2)$$

$$+ (1/8) \cdot U(W + 4) + \cdots$$

where W is the individual's current wealth. If utility took the logarithmic form $U(x) \equiv \ln(x)$ and $W = \$50,000$, for example, the individual's certainty equivalent ξ would only be about $9, even though the gamble has an infinite expected value.

Although it shares the name "utility," $U(\cdot)$ is quite distinct from the ordinal utility function of standard consumer theory. While the latter can be subjected to any monotonic transformation, a von Neumann-Morgenstern utility function is cardinal in that it can only be subjected to transformations of the form $a \cdot U(x) + b$ $(a > 0)$, i.e. transformations which change the origin and/or scale of the vertical axis, but do not affect the "shape" of the function.[3]

To see how this shape determines risk attitudes, consider Figures 1a and 1b. The monotonicity of $U_a(\cdot)$ and $U_b(\cdot)$ in the figures reflects the property of stochastic dominance preference, where one lottery is said to stochastically dominate another one if it can be obtained from it by shifting probability from lower to higher outcome levels.[4] Stochastic dominance preference is thus the probabilistic analogue of the attitude that "more is better."

Consider a gamble offering a $2/3 : 1/3$ chance of the outcomes x' or x''. The points $\bar{x} = (2/3) \cdot x' + (1/3) \cdot x''$ in the figures give the expected value of this

[3] Such transformations are often used to normalize the utility function, for example to set $U(0) = 0$ and $U(M) = 1$ for some large value M.

[4] Thus, for example, a $2/3 : 1/3$ chance of $100 or $20 and a $1/2 : 1/2$ chance of $100 or $30 both stochastically dominate a $1/2 : 1/2$ chance of $100 or $20.

gamble, and $\bar{u}_a = (2/3) \cdot U_a(x') + (1/3) \cdot U_a(x'')$ and $\bar{u}_b = (2/3) \cdot U_b(x') + (1/3) \cdot U_b(x'')$ give its expected utilities for $U_a(\cdot)$ and $U_b(\cdot)$. For the concave utility function $U_a(\cdot)$ we have $U_a(\bar{x}) > \bar{u}_a$, which implies that this individual would prefer a sure gain of $\bar{x}$ (which would yield utility $U_a(\bar{x})$) to the gamble. Since someone with a concave utility function will in fact always prefer receiving the expected value of a gamble to the gamble itself, concave utility functions are termed *risk averse*. For the convex utility function $U_b(\cdot)$ we have $\bar{u}_b > U_b(\bar{x})$, and since this preference for bearing the risk rather than receiving the expected value will also extend to all gambles, $U_b(\cdot)$ is termed *risk loving*. In their famous article, Friedman and Savage (1948) showed how a utility function which was concave at low wealth levels and convex at high wealth levels could explain the behavior of individuals who both incur risk by purchasing lottery tickets as well as avoid risk by purchasing insurance. Algebraically, Arrow (1965) and Pratt (1964) have shown how the degree of concavity of a utility function, as measured by the curvature index $-U''(x)/U'(x)$, determines how risk attitudes, and hence behavior, will vary with wealth or across individuals in a variety of situations. If $U_c(\cdot)$ is at least as risk averse as $U_d(\cdot)$ in the sense that $-U_c''(x)/U_c'(x) \geq -U_d''(x)/U_d'(x)$ for all x, then an individual with utility function $U_c(\cdot)$ would be willing to pay at least as much for insurance against any risk as would someone with utility function $U_d(\cdot)$.

Since a knowledge of $U(\cdot)$ would allow us to predict preferences (and hence behavior) in any risky situation, experimenters and applied decision analysts are frequently interested in eliciting or recovering their subjects' (or clients') von Neumann-Morgenstern utility functions. One method of doing so is termed the *fractile method*. This approach begins by adopting the normalization $U(0) = 0$ and $U(M) = 1$ (see Note 3) and fixing a "mixture probability" $\bar{p}$, say $\bar{p} = 1/2$. The next step involves finding the individual's certainty equivalent ξ_1 of a $1/2:1/2$ chance of M or 0, which implies that $U(\xi_1) = (1/2) \cdot U(M) + (1/2) \cdot U(0) = 1/2$. Finding the certainty equivalents of the $1/2:1/2$ chances of ξ_1 or 0 and of M or ξ_1 yields the values ξ_2 and ξ_3 which solve $U(\xi_2) = 1/4$ and $U(\xi_3) = 3/4$. By repeating this procedure (i.e. $1/8, 3/8, 5/8, 7/8, 1/16, 3/16$, etc.), the utility function can (in the limit) be completely assessed.

Our discussion so far has paralleled the economic literature of the 1960s and 1970s by emphasizing the flexibility of the expected utility model compared to the Pascal-Fermat expected value approach. However, the need to analyze and respond to growing empirical challenges has led economists in the 1980's to concentrate on the behavioral restrictions implied by the expected utility hypothesis. It is to these restrictions that we now turn.

A Modern Perspective: Linearity in the Probabilities as a Testable Hypothesis

As a theory of individual behavior, the expected utility model shares many of the underlying assumptions of standard consumer theory. In each case we assume that the objects of choice, either commodity bundles or lotteries, can be unambiguously and objectively described, and that situations which ultimately imply the same set of availabilities (e.g. the same budget set) will lead to the same choice. In each case we
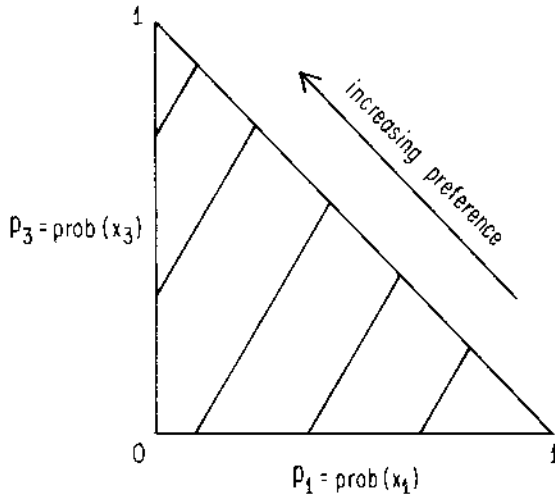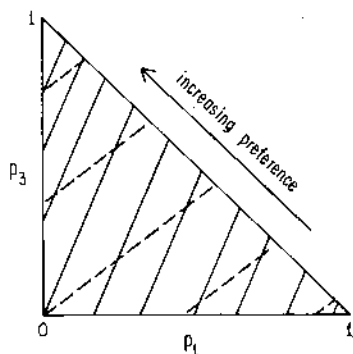
Fig. 2. Expected utility indifference curves in the triangle diagram

also assume that the individual is able to perform the mathematical operations necessary to actually determine the set of availabilities, e.g. to add up the quantities in different sized containers or calculate the probabilities of compound or conditional events. Finally, in each case we assume that preferences are transitive, so that if an individual prefers one object (either a commodity bundle or a risky prospect) to a second, and prefers this second object to a third, he or she will prefer the first object to the third. We shall examine the validity of these assumptions for choice under uncertainty in some of the following sections.

However, the strongest implication of the expected utility hypothesis stems from the form of the expected utility maximand or preference function $\Sigma U(x_i)p_i$. Although this preference function generalizes the expected value form $\Sigma x_i p_i$ by dropping the property of linearity in the payoffs (the x_i's), it retains the other key property of this form, namely linearity in the probabilities.

Graphically, we may illustrate the property of linearity in the probabilities by considering the set of all lotteries or prospects over the fixed outcome levels $x_1 < x_2 < x_3$, which can be represented by the set of all probability triples of the form $P = (p_1, p_2, p_3)$ where $p_i = \text{prob}(x_i)$ and $\Sigma p_i = 1$. Since $p_2 = 1 - p_1 - p_3$, we can represent these lotteries by the points in the unit triangle in the (p_1, p_3) plane, as in Figure 2.[5] Since upward movements in the triangle increase p_3 at the expense of p_2 (i.e. shift probability from the outcome x_2 up to x_3) and leftward movements reduce p_1 to the benefit of p_2 (shift probability from x_1 up to x_2), these movements (and more generally, all northwest movements) lead to stochastically dominating lotteries

[5]Thus if $x_1 = \$20$, $x_2 = \$30$ and $x_3 = \$100$, the prospects in Note 4 would be represented by the points $(p_1, p_3) = (1/3, 2/3)$, $(p_1, p_3) = (0, 1/2)$ and $(p_1, p_3) = (1/2, 1/2)$ respectively. Although it is fair to describe the renewal of interest in this approach as "modern," versions of this diagram go back at least to Marschak (1950).

(solid lines are expected utility indifference curves)
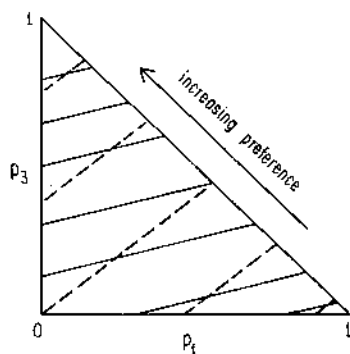(dashed lines are iso-expected value lines)

Fig. 3a. Relatively steep indifference curves of a   Fig. 3b. Relatively flat indifference curves of a risk
risk averter   lover

and would accordingly be preferred. Finally, since the individual's indifference curves in the (p_1, p_3) diagram are given by the solutions to the linear equation

$$(2) \quad \bar{u} = \sum_{i=1}^{3} U(x_i)p_i = U(x_1)p_1 + U(x_2)(1 - p_1 - p_3) + U(x_3)p_3 = \text{constant}$$

they will consist of parallel straight lines of slope $[U(x_2) - U(x_1)]/[U(x_3) - U(x_2)]$, with more preferred indifference curves lying to the northwest. This implies that in order to know an expected utility maximizer's preferences over the entire triangle, it suffices to know the slope of a single indifference curve.

To see how this diagram can be used to illustrate attitudes toward risk, consider Figures 3a and 3b. The dashed lines in the figures are not indifference curves but rather *iso-expected value lines*, i.e. solutions to

$$(3) \quad \bar{x} = \sum_{i=1}^{3} x_i p_i = x_1 p_1 + x_2(1 - p_1 - p_3) + x_3 p_3 = \text{constant}$$

Since northeast movements along these lines do not change the expected value of the prospect but do increase the probabilities of the tail outcomes x_1 and x_3 at the expense of the middle outcome x_2, they are examples of *mean preserving spreads* or "pure" increases in risk (Rothschild and Stiglitz, 1970). When the utility function $U(\cdot)$ is concave (i.e. risk averse), its indifference curves can be shown to be steeper than the iso-expected value lines as in Figure 3a,[6] and such increases in risk will lead to lower indifference curves. When $U(\cdot)$ is convex (risk loving), its indifference curves will be flatter than the iso-expected value lines (as in Figure 3b) and increases in risk will lead

[6]This follows since the slope of the indifference curves is $[U(x_2) - U(x_1)]/[U(x_3) - U(x_2)]$, the slope of the iso-expected value lines is $[x_2 - x_1]/[x_3 - x_2]$, and concavity of $U(\cdot)$ implies $[U(x_2) - U(x_1)]/[x_2 - x_1] > [U(x_3) - U(x_2)]/[x_3 - x_2]$ whenever $x_1 < x_2 < x_3$.

to higher indifference curves. If we compare two different utility functions, the one which is more risk averse (in the above Arrow-Pratt sense) will possess the steeper indifference curves.

Behaviorally, we can view the property of linearity in the probabilities as a restriction on the individual's preferences over probability mixtures of lotteries. If $P^* = (p_1^*, \ldots, p_n^*)$ and $P = (p_1, \ldots, p_n)$ are two lotteries over a common outcome set $\{x_1, \ldots, x_n\}$, the $\alpha:(1 - \alpha)$ probability mixture of P^* and P is the lottery $\alpha P^* + (1 - \alpha)P = (\alpha p_1^* + (1 - \alpha)p_1, \ldots, \alpha p_n^* + (1 - \alpha)p_n)$. This may be thought of as that prospect which yields the same ultimate probabilities over $\{x_1, \ldots, x_n\}$ as the two-stage lottery which offers an $\alpha:(1 - \alpha)$ chance of winning either P^* or P. Since linearity in the probabilities implies that $\Sigma U(x_i)(\alpha p_i^* + (1 - \alpha)p_i) = \alpha \cdot \Sigma U(x_i)p_i^* + (1 - \alpha) \cdot \Sigma U(x_i)p_i$, expected utility maximizers will exhibit the following property, known as the *Independence Axiom* (Samuelson, 1952):

> If the lottery P^* is preferred (resp. indifferent) to the lottery P, then the mixture $\alpha P^* + (1 - \alpha)P^{**}$ will be preferred (resp. indifferent) to the mixture $\alpha P + (1 - \alpha)P^{**}$ for all $\alpha > 0$ and P^{**}.

This property, which is in fact equivalent to linearity in the probabilities, can be interpreted as follows:

> In terms of the ultimate probabilities over the outcomes $\{x_1, \ldots, x_n\}$, choosing between the mixtures $\alpha P^* + (1 - \alpha)P^{**}$ and $\alpha P + (1 - \alpha)P^{**}$ is the same as being offered a coin with a probability of $1 - \alpha$ of landing tails, in which case you will obtain the lottery P^{**}, and being asked before the flip whether you would rather have P^* or P in the event of a head. Now either the coin will land tails, in which case your choice won't have mattered, or else it will land heads, in which case you are 'in effect' back to a choice between P^* or P, and it is only 'rational' to make the same choice as you would before.

Although this is a prescriptive argument, it has played a key role in economists' adoption of expected utility as a descriptive theory of choice under uncertainty. As the evidence against the model mounts, this has lead to a growing tension between those who view economic analysis as the description and prediction of what they consider to be rational behavior and those who view it as the description and prediction of observed behavior. We turn now to this evidence.

Violations of Linearity in the Probabilities

The Allais Paradox and "Fanning Out"

One of the earliest and best known examples of systematic violation of linearity in the probabilities (or equivalently, of the independence axiom) is the well-known *Allais Paradox* (Allais, 1953, 1979). This problem involves obtaining the individual's
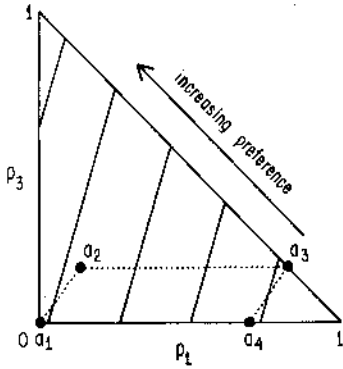
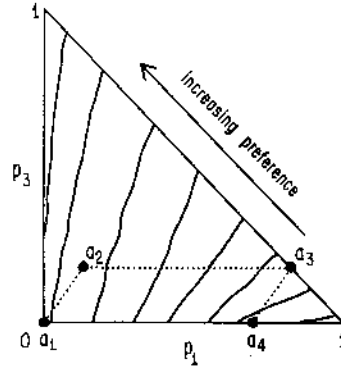Fig. 4a. Expected utility indifference curves and the Allais Paradox

Fig. 4b. Indifference curves which 'fan out' and the Allais Paradox

preferred option from each of the following two pairs of gambles (readers who have never seen this problem may want to circle their own choice from each pair before proceeding):

$$a_1: \left\{ 1.00 \text{ chance of } \$1,000,000 \right. \quad \text{versus} \quad a_2: \left\{ \begin{array}{l} .10 \text{ chance of } \$5,000,000 \\ .89 \text{ chance of } \$1,000,000 \\ .01 \text{ chance of } \$0 \end{array} \right.$$

and

$$a_3: \left\{ \begin{array}{l} .10 \text{ chance of } \$5,000,000 \\ .90 \text{ chance of } \$0 \end{array} \right. \quad \text{versus} \quad a_4: \left\{ \begin{array}{l} .11 \text{ chance of } \$1,000,000 \\ .89 \text{ chance of } \$0 \end{array} \right.$$

Defining $\{x_1, x_2, x_3\} = \{\$0; \$1,000,000; \$5,000,000\}$, these four gambles are seen form a parallelogram in the (p_1, p_3) triangle, as in Figures 4a and 4b. Under the expected utility hypothesis, therefore, a preference for a_1 in the first pair would indicate that the individual's indifference curves were relatively steep (as in Figure 4a), and hence a preference for a_4 in the second pair. In the alternative case of relatively flat indifference curves, the gambles a_2 and a_3 would be preferred.[7] However, researchers such as Allais (1953), Morrison (1967), Raiffa (1968) and Slovic and Tversky (1974) have found that the modal if not majority preferences of subjects has been for a_1 in the first pair and a_3 in the second, which implies that indifference curves are not parallel but rather fan out, as in Figure 4b.

One of the criticisms of this evidence has been that individuals whose choices violated the independence axiom would "correct" themselves once the nature of their violation was revealed by an application of the above coin-flip argument. Thus, while even Savage chose a_1 and a_3 when first presented with this example, he concluded upon reflection that these preferences were in error (Savage, 1954, pp. 101–103).

[7]Algebraically, these cases are equivalent to the expression $[.10 \cdot U(5,000,000) - .11 \cdot U(1,000,000) + .01 \cdot U(0)]$ being negative or positive, respectively.

Although his own reaction was undoubtedly sincere, the hypothesis that individuals would invariably react in such a manner has not been sustained in direct empirical testing. In experiments where subjects were asked to respond to Allais-type problems and then presented with arguments both for and against the expected utility position, neither MacCrimmon (1968), Moskowitz (1974) nor Slovic and Tversky (1974) found predominant net swings toward the expected utility choices.

Additional Evidence of Fanning Out

Although the Allais Paradox was originally dismissed as an isolated example, it is now known to be a special case of a general empirical pattern termed the *common consequence effect*. This effect involves pairs of probability mixtures of the form:

$$b_1: \alpha\delta_x + (1 - \alpha)P^{**} \quad \text{versus} \quad b_2: \alpha P + (1 - \alpha)P^{**}$$

and

$$b_3: \alpha\delta_x + (1 - \alpha)P^* \quad \text{versus} \quad b_4: \alpha P + (1 - \alpha)P^*$$

where δ_x denotes the prospect which yields x with certainty, P involves outcomes both greater and less than x, and P^{**} stochastically dominates P^*.[8] Although the independence axiom clearly implies choices of either b_1 and b_3 (if δ_x is preferred to P) or else b_2 and b_4 (if P is preferred to δ_x), researchers have found a tendency for subjects to choose b_1 in the first pair and b_4 in the second (MacCrimmon, 1968; MacCrimmon and Larsson, 1979; Kahneman and Tversky, 1979; Chew and Waller, 1986). When the distributions δ_x, P, P^* and P^{**} are each over a common outcome set $\{x_1, x_2, x_3\}$, the prospects b_1, b_2, b_3 and b_4 will again form a parallelogram in the (p_1, p_3) triangle, and a choice of b_1 and b_4 again implies indifference curves which fan out, as in Figure 4b.

The intuition behind this phenomenon can be described in terms of the above "coin-flip" scenario. According to the independence axiom, preferences over what would occur in the event of a head should not depend upon what would occur in the event of a tail. In fact, however, they may well depend upon what would otherwise happen.[9] The common consequence effect states that the better off individuals would be in the event of a tail (in the sense of stochastic dominance), the more risk averse they become over what they would receive in the event of a head. Intuitively, if the distribution P^{**} in the pair $\{b_1, b_2\}$ involves very high outcomes, I may prefer not to bear further risk in the unlucky event that I don't receive it, and prefer the sure outcome x over the distribution P in this event (i.e. choose b_1 over b_2). But if P^* in $\{b_3, b_4\}$ involves very low outcomes, I may be more willing to bear risk in the (lucky)

---

[8] The Allais Paradox choices a_1, a_2, a_3, and a_4 correspond to b_1, b_2, b_4 and b_3, where $\alpha = .11$, $x = \$1,000,000$, P is a $10/11 : 1/11$ chance of $\$5,000,000$ or $\$0$, P^* is a sure chance of $\$0$, and P^{**} is a sure chance of $\$1,000,000$. The name of this phenomenon comes from the "common consequence" P^{**} in $\{b_1, b_2\}$ and P^* in $\{b_3, b_4\}$.

[9] As Bell (1985) notes, "winning the top prize of $\$10,000$ in a lottery may leave one much happier than receiving $\$10,000$ as the lowest prize in a lottery."
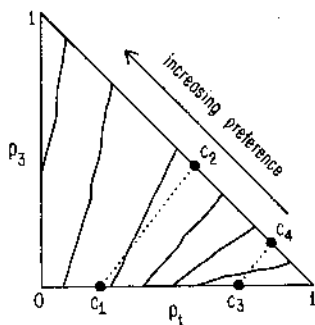
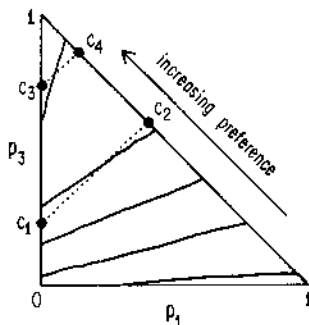Fig. 5a. Indifference curves which fan out and the common ratio effect

Fig. 5b. Indifference curves which fan out and the common ratio effect with negative payoffs

event that I don't receive it, and prefer the lottery P to the outcome x in this case (i.e. choose b_4 over b_3). Note that it is not my beliefs regarding the probabilities in P which are affected here, merely my willingness to bear them.[10]

A second class of systematic violations, stemming from another early example of Allais (1953), is known as the *common ratio effect*. This phenomenon involves pairs of prospects of the form:

$$c_1 : \begin{cases} p \text{ chance of } \$X \\ 1 - p \text{ chance of } \$0 \end{cases} \quad \text{versus} \quad c_2 : \begin{cases} q \text{ chance of } \$Y \\ 1 - q \text{ chance of } \$0, \end{cases}$$

and

$$c_3 : \begin{cases} rp \text{ chance of } \$X \\ 1 - rp \text{ chance of } \$0 \end{cases} \quad \text{versus} \quad c_4 : \begin{cases} rq \text{ chance of } \$Y \\ 1 - rq \text{ chance of } \$0, \end{cases}$$

where $p > q$, $0 < X < Y$ and $r \in (0,1)$, and includes the "certainty effect" of Kahneman and Tversky (1979) and the ingenious "Bergen Paradox" of Hagen (1979) as special cases.[11] Setting $\{x_1, x_2, x_3\} = \{0, X, Y\}$ and plotting these prospects in the (p_1, p_3) triangle, the segments $\overline{c_1 c_2}$ and $\overline{c_3 c_4}$ are seen to be parallel (as in Figure 5a), so that the expected utility model again predicts choices of c_1 and c_3 (if the individual's indifference curves are steep) or else c_2 and c_4 (if they are flat). However, experimental studies have found a systematic tendency for choices to depart from these predictions

---

[10] In a conversation with the author, Kenneth Arrow has offered an alternative phrasing of this argument: The widely maintained hypothesis of decreasing absolute risk aversion asserts that individuals will display more risk aversion in the event of a loss, and less risk aversion in the event of a gain. In the common consequence effect, individuals display more risk aversion in the event of an opportunity loss, and less risk aversion in the event of an opportunity gain.

[11] The former involves setting $p = 1$, and the latter consists of a two-step choice problem where individuals exhibit the effect with $Y = 2X$ and $p = 2q$. The name "common ratio effect" comes from the common value of prob(X)/prob(Y) in the pairs $\{c_1, c_2\}$ and $\{c_3, c_4\}$.
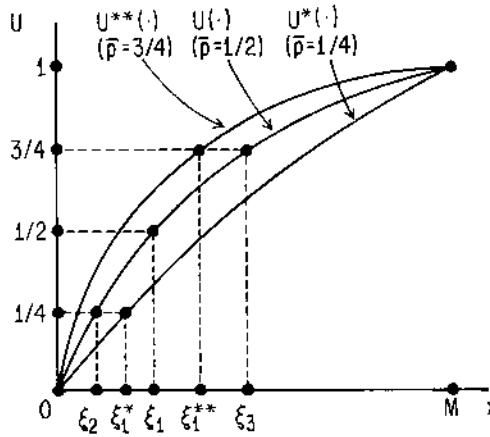
Fig. 6a. 'Recovered' utility functions for mixture probabilities 1/4, 1/2 and 3/4

(~ denotes indifference)
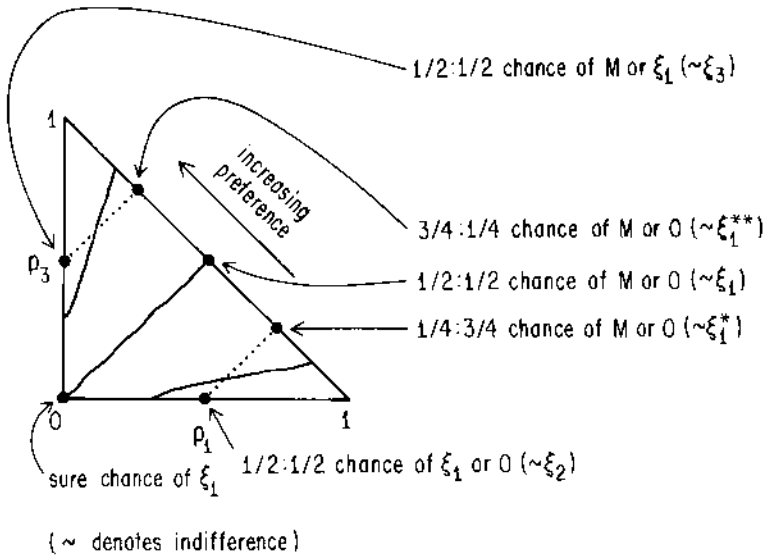
Fig. 6b. Fanning out indifference curves which generate the responses of Figure 6a

and de Neufville (1983, 1984) have found a tendency for higher values of $\bar{p}$ to lead to the "recovery" of higher valued utility functions, as in Figure 6a. By illustrating the gambles used to obtain the values ξ_1, ξ_2 and ξ_3 for $\bar{p} = 1/2$, ξ_1^* for $\bar{p} = 1/4$ and ξ_1^{**} for $\bar{p} = 3/4$, Figure 6b shows that, as with the common consequence and common ratio effects, this *utility evaluation effect* is precisely what would be expected

[12] Kahneman and Tversky (1979), for example, found that 80 percent of their subjects preferred a sure gain of 3,000 Israeli pounds to a .80 chance of winning 4,000, but 65 percent preferred a .20 chance of winning 4,000 to a .25 chance of winning 3,000.

in the direction of preferring c_1 and c_4,[12] which again suggests that indifference curves fan out, as in the figure (Tversky, 1975; MacCrimmon and Larsson, 1979; Chew and Waller, 1986). In a variation on this approach, Kahneman and Tversky (1979) replaced the gains of $\$X$ and $\$Y$ in the above gambles with losses of these magnitudes, and found a tendency to depart from expected utility in the direction of c_2 and c_3. Defining $\{x_1, x_2, x_3\}$ as $\{-Y, -X, 0\}$ (to maintain the condition $x_1 < x_2 < x_3$) and plotting these gambles in Figure 5b, a choice of c_2 and c_3 is again seen to imply that indifference curves fan out. Finally, Battalio, Kagel and MacDonald (1985) found that laboratory rats choosing among gambles which involved substantial variations in their actual daily food intake also exhibited this pattern of choices.

A third class of evidence stems from the elicitation method described in the previous section. In particular, note that there is no reason why the mixture probability $\bar{p}$ must be $1/2$ in this procedure. Picking any other $\bar{p}$ and defining ξ_1^*, ξ_2^* and ξ_3^* as the certainty equivalents of the $\bar{p}: (1 - \bar{p})$ chances of M or 0, ξ_1^* or 0, and M or ξ_1^* yields the equations $U(\xi_1^*) = \bar{p}$, $U(\xi_2^*) = \bar{p}^2$, $U(\xi_3^*) = \bar{p} + (1 - \bar{p})\bar{p}$, etc., and such a procedure can also be used to recover $U(\cdot)$.

Although this procedure should recover the same (normalized) utility function for any mixture probability $\bar{p}$, researchers such as Karmarkar (1974, 1978) and McCord from an individual whose indifference curves departed from expected utility by fanning out.[13]

Non-Expected Utility Models of Preferences

The systematic nature of these departures from linearity in the probabilities have led several researchers to generalize the expected utility model by positing nonlinear functional forms for the individual preference function. Examples of such forms and researchers who have studied them include:

(4) $\qquad \sum \nu(x_i)\pi(p_i)$
$\qquad\qquad$ Edwards (1955)
$\qquad\qquad$ Kahneman and Tversky (1979)

(5) $\qquad \dfrac{\sum \nu(x_i)\pi(p_i)}{\sum \pi(p_i)}$
$\qquad\qquad$ Karmarkar (1978)

(6) $\qquad \dfrac{\sum \nu(x_i)p_i}{\sum \tau(x_i)p_i}$
$\qquad\qquad$ Chew (1983)
$\qquad\qquad$ Fishburn (1983)

(7) $\sum \nu(x_i)[g(p_1 + \cdots + p_i) - g(p_1 + \cdots + p_{i-1})]$ Quiggin (1982)

(8) $\qquad \sum \nu(x_i)p_i + [\sum \tau(x_i)p_i]^2$
$\qquad\qquad$ Machina (1982)

[13]Having found that ξ_1 which solves $U(\xi_1) = (1/2) \cdot U(M) + (1/2) \cdot U(0)$, choose $\{x_1, x_2, x_3\} = \{0, \xi_1, M\}$, so that the indifference curve through $(0,0)$ (i.e. a sure gain of ξ_1) also passes through $(1/2, 1/2)$ (a $1/2 : 1/2$ chance of M or 0). The order of $\xi_1, \xi_2, \xi_3, \xi_1^*$ and ξ_1^{**} in Figure 6a is derived from the individual's preference ordering over the five distributions in Figure 6b for which they are the respective certainty equivalents.

Many (though not all) of these forms are flexible enough to exhibit the properties of stochastic dominance preference, risk aversion/risk preference and fanning out, and (6) and (7) have proven to be particularly useful both theoretically and empirically. Additional analyses of these forms can be found in Chew, Karni and Safra (1987), Fishburn (1964), Segal (1984) and Yaari (1987).

Although such forms allow for the modelling of preferences which are more general than those allowed by the expected utility hypothesis, each requires a different set of conditions on its component functions $\nu(\cdot)$, $\pi(\cdot)$, $\tau(\cdot)$ or $g(\cdot)$ for the properties of stochastic dominance preference, risk aversion/risk preference, comparative risk aversion, etc. In particular, the standard expected utility results linking properties of the function $U(\cdot)$ to such aspects of behavior will generally not extend to the corresponding properties of the function $\nu(\cdot)$ in the above forms. Does this mean that the study of non-expected utility preferences requires us to abandon the vast body of theoretical results and intuition we have developed within the expected utility framework?

Fortunately, the answer is no. An alternative approach to the analysis of non-expected utility preferences proceeds not by adopting a specific nonlinear function, but rather by considering nonlinear functions in general, and using calculus to extend the results from expected utility theory in the same manner in which it is typically used to extend results involving linear functions.[14]

Specifically, consider the set of all probability distributions $P = (p_1, \ldots, p_n)$ over a fixed outcome set $\{x_1, \ldots, x_n\}$, so that the expected utility preference function can be written as $V(P) = V(p_1, \ldots, p_n) \equiv \Sigma U(x_i)p_i$, and think of $U(x_i)$ not as a "utility level" but rather as the coefficient of $p_i = \text{prob}(x_i)$ in this linear function. If we plot these coefficients against x_i as in Figure 7, the expected utility results of the previous section can be stated as:

Stochastic Dominance Preference: $V(\cdot)$ will exhibit global stochastic dominance preference if and only if the coefficients $\{U(x_i)\}$ are increasing in x_i, as in the figure.

Risk Aversion: $V(\cdot)$ will exhibit global risk aversion if and only if the coefficients $\{U(x_i)\}$ are concave in x_i,[15] as in the figure.

Comparative Risk Aversion: The expected utility preference function $V^*(P) \equiv \Sigma U^*(x_i)p_i$ will be at least as risk averse as $V(\cdot)$ if and only if the coefficients $\{U^*(x_i)\}$ are at least as concave in x_i as $\{U(x_i)\}$.[16]

Now take the case where the individual's preference function $\mathscr{V}(P) = \mathscr{V}(p_1, \ldots, p_n)$ is not linear (i.e. not expected utility) but at least differentiable, and consider its partial derivatives $\mathscr{U}(x_i; P) \equiv \partial \mathscr{V}(P)/\partial p_i = \partial \mathscr{V}(P)/\partial \, \text{prob}(x_i)$. Pick some probability distribution P_0 and plot these $\mathscr{U}(x_i; P_0)$ values against x_i. If they

---

[14] Readers who wish to skip the details of this approach may proceed to the next section.

[15] As in Note 6, this is equivalent to the condition that $[U(x_{i+1}) - U(x_i)]/[x_{i+1} - x_i] < [U(x_i) - U(x_{i-1})]/[x_i - x_{i-1}]$ for all i.

[16] This is equivalent to the condition that $U^*(x_i) \equiv \rho(U(x_i))$ for some increasing concave function $\rho(\cdot)$.
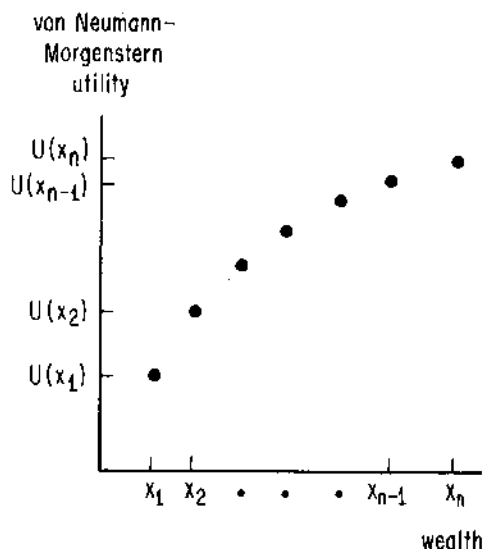
Fig. 7. von Neumann-Morgenstern utilities as coefficients of the expected utility preference function $V(p_1, \ldots, p_n) \equiv \Sigma U(x_i)p_i$

are again increasing in x_i, it is clear that any infinitesimal stochastically dominating shift from P_0, such as a decrease in some p_i and matching increase in p_{i+1}, will be preferred. If they are again concave in x_i, any infinitesimal mean preserving spread, such as a drop in p_i and (mean preserving) rise in p_{i-1} and p_{i+1}, will make the individual worse off. In light of this correspondence between the coefficients $\{U(x_i)\}$ of an expected utility preference function $V(\cdot)$ and the partial derivatives $\{\mathcal{U}(x_i; P_0)\}$ of the non-expected utility preference function $\mathcal{V}(\cdot)$, we refer to $\{\mathcal{U}(x_i; P_0)\}$ as the individual's local utility indices at P_0.

Of course, the above results will only hold precisely for infinitesimal shifts from the distribution P_0. However, we can exploit another result from standard calculus to show how "expected utility" results may be applied to the exact global analysis of non-expected utility preferences. Recall that in many cases, a differentiable function will exhibit a specific global property if and only if that property is exhibited by its linear approximations at each point. For example, a differentiable function will be globally nondecreasing if and only if its linear approximations are non-decreasing at each point. In fact, most of the fundamental properties of risk attitudes and their expected utility characterizations are precisely of this type. In particular, it can be shown that:

Stochastic Dominance Preference: A non-expected utility preference function $\mathcal{V}(\cdot)$ will exhibit global stochastic dominance preference if and only if its local utility indices $\{\mathcal{U}(x_i; P)\}$ are increasing in x_i at each distribution P.

Risk Aversion: $\mathcal{V}(\cdot)$ will exhibit global risk aversion if and only if its local utility indices $\{\mathcal{U}(x_i; P)\}$ are concave in x_i at each distribution P.
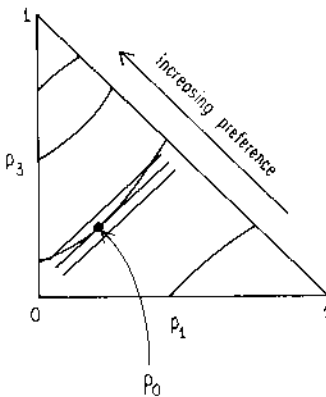
(solid lines are local expected utility approximation
to non-expected utility indifference curves at P_0)
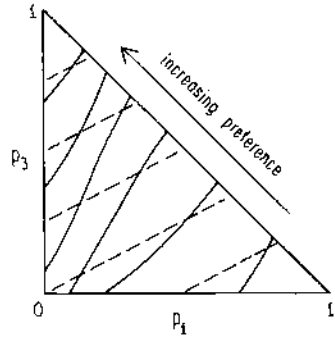
(dashed lines are iso-expected value lines)

Fig. 8a. Tangent 'expected utility' approximation to non-expected utility indifference curves

Fig. 8b. Risk aversion of every local expected utility approximation is equivalent to global risk aversion

Comparative Risk Aversion: The preference function $\mathscr{V}^*(\cdot)$ will be globally at least as risk averse as $\mathscr{V}(\cdot)$[17] if and only if its local utility indices $\{\mathscr{U}^*(x_i; P)\}$ are at least as concave in x_i as $\{\mathscr{U}(x_i; P)\}$ at each P.

Figures 8a and 8b give a graphical illustration of this approach for the outcome set $\{x_1, x_2, x_3\}$. Here the solid curves denote the indifference curves of the non-expected utility preference function $\mathscr{V}(P)$. The parallel lines near the lottery P_0 denote the tangent "expected utility" indifference curves that correspond to the local utility indices $\{\mathscr{U}(x_i; P_0)\}$ at P_0. As always with differentiable functions, an infinitesimal change in the probabilities at P_0 will be preferred if and only if they would be preferred by this tangent linear (i.e. expected utility) approximation. Figure 8b illustrates the above "risk aversion" result: It is clear that these indifference curves will be globally risk averse (averse to mean preserving spreads) if and only if they are everywhere steeper than the dashed iso-expected value lines. However, this is equivalent to all of their *tangents* being steeper than these lines, which is in turn equivalent to all of their local expected utility approximations being risk averse, or in other words, to the local utility indices $\{\mathscr{U}(x_i; P)\}$ being concave in x_i at each distribution P.

My fellow researchers and I have shown how this and similar techniques can be applied to further extend the results of expected utility theory to the case of non-expected utility preferences, to characterize and explore the implications of preferences which "fan out," and to conduct new and more general analyses of economic behavior under uncertainty (Machina, 1982; Chew, 1983; Fishburn, 1984; Epstein, 1985; Allen, 1987; Chew, Karni and Safra, 1987). However, while I feel that

---

[17] For the appropriate generalizations of the expected utility concepts of "at least as risk averse" in this context, see Machina (1982, 1984).

they constitute a useful and promising response to the phenomenon of non-linearities in the probabilities, these models do not provide solutions to the more problematic empirical phenomena of the following sections.

The Preference Reversal Phenomenon

The Evidence

The finding now known as the preference reversal phenomenon was first reported by psychologists Lichtenstein and Slovic (1971). In this study, subjects were first presented with a number of pairs of bets and asked to choose one bet out of each pair. Each of these pairs took the following form:

$$P\text{-bet}: \begin{cases} p \text{ chance of } \$X \\ 1 - p \text{ chance of } \$x \end{cases} \text{ versus } \$\text{-bet}: \begin{cases} q \text{ chance of } \$Y \\ 1 - q \text{ chance of } \$y, \end{cases}$$

where X and Y are respectively greater than x and y, p is greater than q, and Y is greater than X (the names "P-bet" and "$\$$-bet" come from the greater probability of winning in the first bet and greater possible gain in the second). In some cases, x and y took on small negative values. The subjects were next asked to "value" (state certainty equivalents for) each of these bets. The different valuation methods used consisted of (1) asking subjects to state their minimum selling price for each bet if they were to own it, (2) asking them to state their maximum bid price for each bet if they were to buy it, and (3) the elicitation procedure of Becker, DeGroot and Marschak (1964), in which it is in a subject's best interest to reveal his or her true certainty equivalents.[18] In the latter case, real money was used.

The expected utility model, as well as each of the non-expected utility models of the previous section, clearly implies that the bet which is actually chosen out of each pair will also be the one which is assigned the higher certainty equivalent.[19] However, Lichtenstein and Slovic found a systematic tendency for subjects to violate this prediction by choosing the P-bet in a direct choice but assigning a higher value to the $\$$-bet. In one experiment, for example, 127 out of 173 subjects assigned a higher sell price to the $\$$-bet in every pair in which the P-bet was chosen. Similar findings were obtained by Lindman (1971), and in an interesting variation on the usual experimental setting, by Lichtenstein and Slovic (1973) in a Las Vegas casino where customers actually staked (and hence sometimes lost) their own money. In another real-money

---

[18] Roughly speaking, the subject states a value for the item, and then the experimenter draws a random price. If the price is above the stated value, the subject forgoes the item and receives the price. If the drawn price is below the state value, the subject keeps the item. The reader can verify that under such a scheme it can never be in a subject's best interest to report anything other than his or her true value.

[19] Economic theory tells us that income effects could cause an individual to assign a lower bid price to the object which, if both were free, would actually be preferred. However, this reversal should not occur for either selling prices or the Becker, DeGroot and Marschak elicitations. For evidence on sell price/bid price disparities, see Knetsch and Sinden (1984) and the references cited there.

experiment, Mowen and Gentry (1980) found that groups who could discuss their (joint) decisions were, if anything, more likely than individuals to exhibit the phenomenon.

Although the above studies involved deliberate variations in design in order to check for the robustness of this phenomenon, they were nevertheless received skeptically by economists, who perhaps not unnaturally felt they had more at stake than psychologists in this type of finding. In an admitted attempt to "discredit" this work, economists Grether and Plott (1979) designed a pair of experiments which, by correcting for issues of incentives, income effects, strategic considerations, ability to indicate indifference and other items, would presumably not generate this phenomenon. They nonetheless found it in both experiments. Further design modifications by Pommerehne, Schneider and Zweifel (1982) and Reilly (1982) yielded the same results. Finally, the phenomenon has been found to persist (although in mitigated form) even when subjects are allowed to engage in experimental market transactions involving the gambles (Knez and Smith, 1986), or when the experimenter is able to act as an arbitrageur and make money off of such reversals (Berg, Dickhaut and O'Brien, 1983).

Two Interpretations of this Phenomenon

How you interpret these findings depends on whether you adopt the worldview of an economist or a psychologist. An economist would reason as follows: Each individual possesses a well-defined preference relation over objects (in this case lotteries), and information about this relation can be gleaned from either direct choice questions or (properly designed) valuation questions. Someone exhibiting the preference reversal phenomenon is therefore telling us that he or she (1) is indifferent between the P-bet and some sure amount ξ_P, (2) strictly prefers the P-bet to the \$-bet, and (3) is indifferent between the \$-bet and an amount $\xi_\$$ greater than ξ_P. Assuming they prefer $\xi_\$$ to the lesser amount ξ_P, this implies that their preferences over these four objects are cyclic or intransitive.

Psychologists on the other hand would deny the premise of a common underlying mechanism generating both choice and valuation behavior. Rather, they view choice and valuation (even different forms of valuation) as distinct processes, subject to possibly different influences. In other words, individuals exhibit what are termed *response mode effects*. Excellent discussions and empirical examinations of this phenomenon and its implications for the elicitation of probabilistic beliefs and utility functions can be found in Hogarth (1975), Slovic, Fischhoff and Lichtenstein (1982), Hershey and Schoemaker (1985) and MacCrimmon and Wehrung (1986). In reporting how the response mode study of Slovic and Lichtenstein (1968) led them to actually predict the preference reversal phenomenon, I can do no better than quote the authors themselves:

"The impetus for this study [Lichtenstein and Slovic (1971)] was our observation in our earlier 1968 article that choices among pairs of gambles appeared to

be influenced primarily by probabilities of winning and losing, whereas buying and selling prices were primarily determined by the dollar amounts that could be won or lost. ... In our 1971 article, we argued that, if the information in a gamble is processed differently when making choices and setting prices, it should be possible to construct pairs of gambles such that people would choose one member of the pair but set a higher price on the other."[11]

<div align="right">Slovic and Lichtenstein (1983)</div>

Implications of the Economic Worldview

The issue of intransitivity is new neither to economics nor to choice under uncertainty. May (1954), for example, observed intransitivities in pairwise rankings of three alternative marriage partners, where each candidate was rated highly in two of three attributes (intelligence, looks, wealth) and low in the third. In an uncertain context, Blyth (1972) has adapted this approach to construct a set of random variables $(\tilde{x}, \tilde{y}, \tilde{z})$ such that $\text{prob}(\tilde{x} > \tilde{y}) = \text{prob}(\tilde{y} > \tilde{z}) = \text{prob}(\tilde{z} > \tilde{x}) = 2/3$, so that individuals making pairwise choices on the basis of these probabilities would also be intransitive. In addition to the preference reversal phenomenon, Edwards (1954, pp. 404–405) and Tversky (1969) have also observed intransitivities in preferences over risky prospects. On the other hand, researchers have shown that many aspects of economic theory, in particular the existence of demand functions and of general equilibrium, are surprisingly robust to dropping the assumption of transitivity (Sonnenschein, 1971; Mas-Colell, 1974; Shafer, 1974).

In any event, economists have begun to develop and analyze models of nontransitive preferences over lotteries. The leading example of this is the "expected regret" model developed independently by Bell (1982), Fishburn (1982) and Loomes and Sugden (1982). In this model of pairwise choice, the von Neumann-Morgenstern utility function $U(x)$ is replaced by a *regret/rejoice function* $r(x, y)$ which represents the level of satisfaction (or if negative, dissatisfaction) the individual would experience if he or she were to receive the outcome x when the alternative choice would have yielded the outcome y (this function is assumed to satisfy $r(x, y) \equiv -r(y, x)$). In choosing between statistically independent gambles $P^* = (p_1^*, \ldots, p_n^*)$ and $P = (p_1, \ldots, p_n)$ over a common outcome set $\{x_1, \ldots, x_n\}$, the individual will choose P^* if the expectation $\sum_i \sum_j r(x_i, x_j) p_i^* p_j$ is positive, and P if it is negative.

Note that when the regret/rejoice function takes the special form $r(x, y) \equiv U(x) - U(y)$ this model reduces to the expected utility model, since we have

$$(9) \qquad \sum_i \sum_j r(x_i, x_j) p_i^* p_j$$

$$= \sum_i \sum_j \left[U(x_i) - U(x_j) \right] p_i^* p_j \equiv \sum_i U(x_i) p_i^* - \sum_j U(x_j) p_j$$
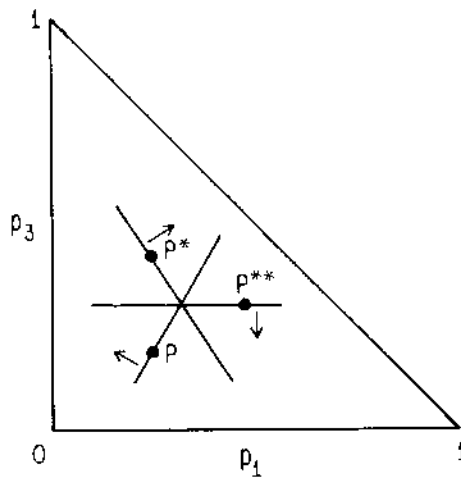
Fig. 9. 'Indifference curves' for the expected regret model

so that the individual will prefer P^* to P if and only if $\sum_i U(x_i)p_i^* > \sum_j U(x_j)p_j$.[20] However, in general such an individual will neither be an expected utility maximizer nor have transitive preferences.

However, this intransitivity does not prevent us from graphing such preferences, or even applying "expected utility" analysis to them. To see the former, consider the case when the individual is facing alternative independent lotteries over a common outcome set $\{x_1, x_2, x_3\}$, so that we may again use the triangle diagram to illustrate their "indifference curves," which will appear as in Figure 9. In such a case it is important to understand what is and is not still true of these indifference curves. The curve through P will still correspond to the set of lotteries that are indifferent to P, and it will still divide the set of lotteries that are strictly preferred to P (the points in the direction of the arrow) from the ones to which P is strictly preferred. Furthermore, if (as in the figure) P^* lies above the indifference curve through P, then P will lie below the indifference curve through P^* (i.e. the individual's ranking of P and P^* will be unambiguous). However, unlike indifference curves for transitive preferences, these curves will cross,[21] and preferences over the lotteries P, P^* and P^{**} are seen to form an intransitive cycle. But in regions where the indifference curves do not cross (such as near the origin) the individual will be indistinguishable from someone with transitive (albeit non-expected utility) preferences.

To see how expected utility results can be extended to this nontransitive framework, fix a lottery $P = (p_1, \ldots, p_n)$ and consider the question of when an (independent) lottery $P^* = (p_1^*, \ldots, p_n^*)$ will be preferred or not preferred to P. Since

[20] When $r(x, y)$ takes the form $r(x, y) \equiv v(x)\tau(y) - v(y)\tau(x)$, this model will reduce to the (transitive) model of equation (6). This is the most general form of the model which is compatible with transitivity.
[21] In this model the indifference curves will all cross at the same point. This point will thus be indifferent to all lotteries in the triangle.

$r(x, y) \equiv -r(y, x)$ implies $\sum_i \sum_j r(x_i, x_j) p_i p_j \equiv 0$, we have that P^* will be preferred to P if and only if

$$(10) \quad 0 < \sum_i \sum_j r(x_i, x_j) p_i^* p_j = \sum_i \sum_j r(x_i, x_j) p_i^* p_j - \sum_i \sum_j r(x_i, x_j) p_i p_j$$

$$= \sum_i \left[\sum_j r(x_i, x_j) p_j \right] p_i^* - \sum_i \left[\sum_j r(x_i, x_j) p_j \right] p_i$$

$$= \sum_i \phi(x_i; P) p_i^* - \sum_i \phi(x_i; P) p_i$$

In other words, P^* will be preferred to P if and only if it implies a higher expectation of the "utility function" $\phi(x_i; P) \equiv \sum_j r(x_i, x_j) p_j$ than P. Thus if $\phi(x_i; P)$ is increasing in x_i for all lotteries P the individual will exhibit global stochastic dominance preference, and if $\phi(x_i; P)$ is concave in x_i for all P the individual will exhibit global risk aversion, even though he or she is not necessarily transitive (these conditions will clearly be satisfied if $r(x, y)$ is increasing and concave in x). The analytics of expected utility theory are robust indeed.

The developers of this model have shown how specific assumptions on the form of the regret/rejoice function will generate the common consequence effect, the common ratio effect, the preference reversal phenomenon, and other observed properties of choice over lotteries. The theoretical and empirical prospects for this approach accordingly seem quite impressive.

Implications of the Psychological Worldview

On the other hand, how should economists respond if it turns out that the psychologists are right, and the preference reversal phenomenon really is generated by some form of response mode effect (or effects)? In that case, the first thing to do would be to try to determine if there were analogues of such effects in real-world economic situations.[22] Will individuals behave differently when determining their valuation of an object (e.g. reservation bid on a used car) than when reacting to a fixed and non-negotiable price for the same object? Since a proper test of this would require correcting for any possible strategic and/or information-theoretic (e.g. signalling) issues, it would not be a simple undertaking. However, in light of the experimental evidence, I feel it is crucial that we attempt it.

Say we found that response mode effects did not occur outside of the laboratory. In that case we could rest more easily, although we could not forget about such issues completely: experimenters testing other economic theories and models (e.g. auctions) would have to be forever mindful of the possible influence of the particular response mode used in their experimental design.

---

[22] It is important to note that neither the evidence of response mode effects (e.g. Slovic, 1975) nor their implications for economic analysis are confined to the case of choice under uncertainty.

On the other hand, what if we did find response mode effects out in the field? In that case we would want to determine, perhaps by going back to the laboratory, whether the rest of economic theory remained valid provided the response mode is held constant. If this were true, then with further evidence on exactly how the response mode mattered, we could presumably incorporate it as a new independent variable into existing theories. Since response modes tend to be constant within a given economic model, e.g. quantity responses to fixed prices in competitive markets, valuation announcements (truthful or otherwise) in auctions, etc., we should expect most of the testable implications of this approach to appear as cross-institutional predictions, such as systematic violations of the various equivalency results involving prices versus quantities or second price-sealed bid versus oral English auctions. In such a case, the new results and insights regarding our theories of institutions and mechanisms could be exciting indeed.[23]

Framing Effects

Evidence

In addition to response mode effects, psychologists have uncovered an even more disturbing phenomenon, namely that alternative means of representing or "framing" probabilistically equivalent choice problems will lead to systematic differences in choice. An early example of this phenomenon was reported by Slovic (1969), who found that offering a gain or loss contingent on the joint occurrence of four independent events with probability p elicited different responses than offering it on the occurrence of a single event with probability p^4 (all probabilities were stated explicitly). In comparison with the single-event case, making a gain contingent on the joint occurrence of events was found to make it more attractive, and making a loss contingent on the joint occurrence of events made it more unattractive.

In another study, Payne and Braunstein (1971) used pairs of gambles of the type illustrated in Figure 10. Each of the gambles in the figure, known as a *duplex gamble*, involves spinning the pointers on both its "gain wheel" (on the left) and its "loss wheel" (on the right), with the individual receiving the sum of the resulting amounts. Thus an individual choosing Gamble A would win $.40 with probability .3 (i.e. if the pointer in the gain wheel landed up and the pointer in the loss wheel landed down), would lose $.40 with probability .2 (if the pointers landed in the opposite positions), and would break even with probability .5 (if the pointers landed either both up or both down). An examination of Gamble B reveals that it has an identical underlying distribution, so that subjects should be indifferent between the two gambles regardless

[23]A final "twist" on the preference reversal phenomenon: Holt (1986) and Karni and Safra (1987) have shown how the procedures used in most of these studies will only lead to truthful revelation of preferences under the added assumption that the individual satisfies the independence axiom, and has given examples of transitive non-expected utility preference rankings which lead to the typical "preference reversal" choices. How (and whether) experimenters will be able to address this issue remains to be seen.

$.40   -$.40

.6   .5

.4   .5

$0   GAMBLE A $0

$.40   -$.40

.5   .4

.5   .6

$0   GAMBLE B $0
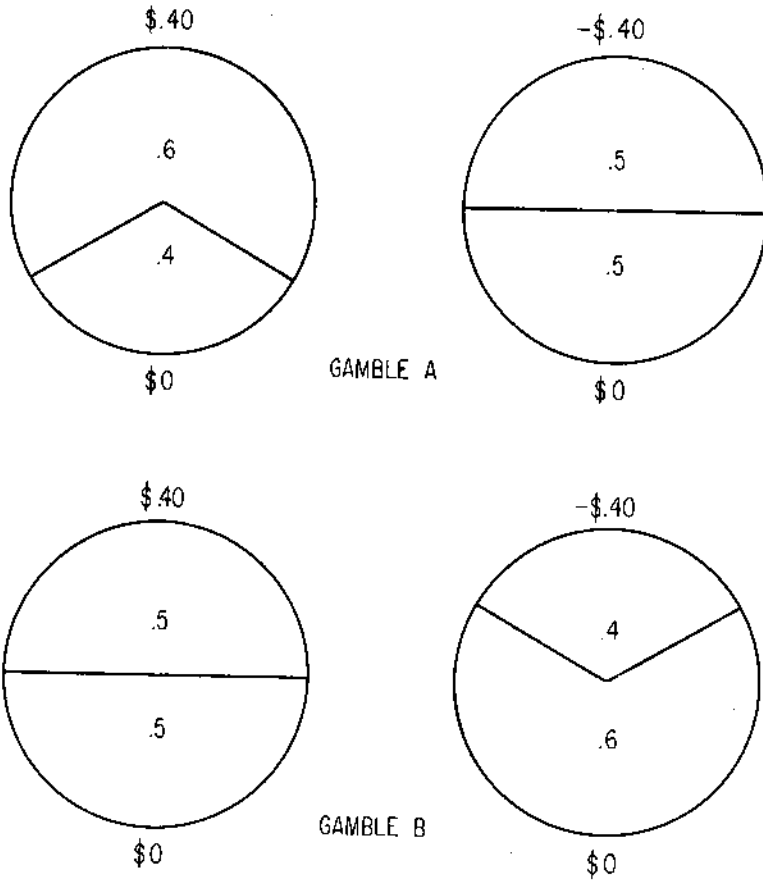
Fig. 10. Duplex gambles with identical underlying distributions

of their risk preferences. However, Payne and Braunstein found that individuals in fact chose between such gambles (and indicated nontrivial strengths of preference) in manners which were systematically affected by the attributes of the component wheels. When the probability of winning in the gain wheel was greater than the probability of losing in the loss wheel for each gamble (as in the figure), subjects tended to choose the gamble whose gain wheel yielded the greater probability of a gain (Gamble A). In cases where the probabilities of losing in the loss wheels were respectively greater than the probabilities of winning in the gain wheels, subjects tended toward the gamble with the lower probability of losing in the loss wheel.

Finally, although the gambles in Figure 10 possess identical underlying distributions, continuity suggests that a slight worsening of the terms of the preferred gamble could result in a pair of non-equivalent duplex gambles in which the individual will actually choose the one with the stochastically dominated underlying distribution. In an experiment where the subjects were allowed to construct their own duplex gambles by choosing one from a pair of prospects involving gains and one from a pair of

prospects involving losses, stochastically dominated prospects were indeed chosen (Tversky and Kahneman, 1981).[24]

A second class of framing effects involves the phenomenon of a *reference point*. Theoretically, the variable which enters an individual's von Neumann-Morgenstern utility function should be total (i.e. final) wealth, and gambles phrased in terms of gains and losses should be combined with current wealth and re-expressed as distributions over final wealth levels before being evaluated. However, economists since Markowitz (1952) have observed that risk attitudes over gains and losses are more stable than can be explained by a fixed utility function over final wealth, and have suggested that the utility function might be best defined in terms of changes from the "reference point" of current wealth. This stability of risk attitudes in the face of wealth variations has also been observed in several experimental studies.[25]

Markowitz (p. 155) also suggested that certain circumstances may cause the individual's reference point to temporarily deviate from current wealth. If these circumstances include the manner in which a given problem is verbally described, then differing risk attitudes over gains and losses can lead to different choices depending upon the exact description. A simple example of this, from Kahneman and Tversky (1979), involves the following two questions:

> In addition to whatever you own, you have been given 1,000 (Israeli pounds). You are now asked to choose between a 1/2 : 1/2 chance of a gain of 1,000 or 0 or a sure gain of 500.

> and

> In addition to whatever you own, you have been given 2,000. You are now asked to choose between a 1/2 : 1/2 chance of loss of 1,000 or 0 or a sure loss of 500.

These two problems involve identical distributions over final wealth. However, when put to two different groups of subjects, 84 percent chose the sure gain in the first problem but 69 percent chose the 1/2 : 1/2 gamble in the second. A nonmonetary version of this type of example, from Tversky and Kahneman (1981, 1986), posits the following scenario:

> 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

---

[24] Subjects were asked to choose either (A) a sure gain of $240 or (B) a 1/4 : 3/4 chance of $1,000 or $0, and to choose either (C) a sure loss of $750 or (D) a 3/4 : 1/4 chance of −$1,000 or 0. 84 percent chose A over B and 87 percent chose D over C, even though B + C dominates A + D, and choices over the combined distributions were unanimous when they were presented explicitly.

[25] See the discussion and references in Machina (1982, pp. 285–86).

consequences of the programs are as follows:

If program A is adopted, 200 people will be saved.

If Program B is adopted, there is 1/3 probability that 600 people will be saved, and 2/3 probability that no people will be saved.

Seventy-two percent of the subjects who were presented with this form of the question chose Program A. A second group was given the same initial information, but the descriptions of the programs were changed to read:

If Program C is adopted 400 people will die.

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

Although this statement of the problem is once again identical to the former one, 78 percent of the respondents chose Program D.

In other studies, Schoemaker and Kunreuther (1979), Hershey and Schoemaker (1980), McNeil, Pauker, Sox and Tversky (1982) and Slovic, Fischhoff and Lichtenstein (1982) have found that subjects' choices in otherwise identical problems will depend upon whether they are phrased as decisions whether or not to gamble or whether or not to insure, whether the statistical information for different therapies is presented in terms of cumulative survival probabilities or cumulative mortality probabilities, etc. For similar examples of this phenomenon in non-stochastic situations, see Thaler (1980).

In a final class of examples, not based on reference point effects, Moskowitz (1974) and Keller (1982) found that the proportion of subjects choosing in conformance with or in violation of the independence axiom in examples like the Allais Paradox was significantly affected by whether the problems were described in the standard matrix form (e.g. Raiffa, 1968, p. 7), decision tree form, or as minimally structured written statements. Interestingly enough, the form which was judged the "clearest representation" by the majority of Moskowitz's subjects (the tree form) led to the lowest degree of consistency with the independence axiom, the highest proportion of fanning out choices, and the highest persistency rate of these choices (pp. 234, 237–38).

Two Issues Regarding Framing

The replicability and pervasiveness of the above types of examples is indisputable. However, before being able to assess their implications for economic modelling we need to resolve two issues.

The first issue is whether these experimental observations possess any analogue outside of the laboratory. Since real-world decision problems do not present themselves as neatly packaged as the ones on experimental questionnaires, monitoring such effects would not be as straightforward. However this does not mean that they do not exist, or that they cannot be objectively observed or quantitatively measured. The real-world example which comes most quickly to mind, and is presumably of no small

importance to the involved parties, is whether gasoline price differentials should be represented as "cash discounts" or "credit surcharges." Similarly, Russo, Krieser and Miyashita (1975) and Russo (1977) found that the practice, and even method, of displaying unit price information in supermarkets (information which consumers could calculate for themselves) affected both the level and distribution of consumer expenditures. The empirical marketing literature is no doubt replete with findings that we could legitimately interpret as real-world framing effects.

The second, more difficult issue is that of the independent observability of the particular frame that an individual will adopt in a given problem. In the duplex gamble and matrix/decision tree/written statement examples of the previous section, the different frames seem unambiguously determined by the form of presentation. However, in instances where framing involves the choice of a reference point, which presumably include the majority of real-world cases, this point might not be objectively determined by the form of presentation, and might be chosen differently, and what is worse, unobservably, by each individual.[26] In a particularly thorough and insightful study, Fischhoff (1983) presented subjects with a written decision problem which allowed for different choices of a reference point, and explored different ways of predicting which frame individuals would adopt, in order to be able to predict their actual choices. While the majority choice of subjects was consistent with what would appear to be the most appropriate frame, Fischhoff noted "the absence of any relation within those studies between [separately elicited] frame preference and option preference." Indeed to the extent that frame preferences varied across his experiments, they did so inversely to the incidence of the predicted choice.[27] If such problems can occur in predicting responses to specific written questions in the laboratory, imagine how they could plague the modelling of real world choice behavior.

Framing Effects and Economic Analysis: Have We Already Solved this Problem?

How should we respond if it turns out that framing actually is a real-world phenomenon of economic relevance, and in particular, if individuals' frames cannot always be observed? I would argue that the means of responding to this issue can already be found in the "tool box" of existing economic analysis.

Consider first the case where the frame of a given economic decision problem, even though it should not matter from the point of view of standard theory, can at least be independently and objectively observed. I believe that economists have in fact already solved such a problem in their treatment of the phenomenon of "uninformative advertising." Although it is hard to give a formal definition of this term, it is widely felt that economic theory is hard put to explain a large proportion of current advertising in terms of traditional informational considerations.[28] However, this has

---

[26] This is not to say that well-defined reference points never exist. The reference points involved in credit surcharges vs. cash discounts, for example, seem unambiguous.

[27] Fischhoff (1983, pp. 115–16). Fischhoff notes that "If one can only infer frames from preferences after assuming the truth of the theory, one runs the risk of making the theory itself untestable."

[28] A wonderful example, offered by my colleague Joel Sobel, are milk ads which make no reference to either price or a specific dairy. What could be a more well-known commodity than milk?

hardly led economists to abandon classical consumer theory. Rather, models of uninformative advertising proceed by quantifying this variable (e.g. air time) and treating it as an additional independent variable in the utility and/or demand function. Standard results like the Slutsky equation need not be abandoned, but rather simply reinterpreted as properties of demand functions holding this new variable constant. The amount of advertising itself is determined as a maximizing variable on the part of the firm (given some cost curve), and can be subjected to standard comparative static analysis.

In the case when decision frames can be observed, framing effects can presumably be modelled in an analogous manner. To do so, we would begin by adopting a method of quantifying, or at least categorizing, frames. The second step, some of which has of course already been done, is to study both the effect of this new independent variable holding the standard economic variables constant, and conversely, to retest our standard economic theories in conditions where we carefully held the frame fixed. With any luck we would find that, holding the frame constant, the Slutsky equation still held.

The next step in any given modelling situation would be to ask "who determines the frame?" If (as with advertising) it is the firm, then the effect of the frame upon consumer demand, and hence upon firm profits, can be incorporated into the firm's maximization problem, and the choice of the frame as well as the other relevant variables (e.g. prices and quantities) can be simultaneously determined and subjected to comparative static analysis, just as in the case of uninformative advertising.

A seemingly more difficult case is when the individual chooses the frame (for example, a reference point) and this choice cannot be observed. Although we should not forget the findings of Fischhoff (1983), assume that this choice is at least systematic in the sense that the consumer will jointly choose the frame and make the subsequent decision in a manner which maximizes a "utility function" which depends both on the decision and the choice of frame. In other words, individuals make their choices as part of a joint maximization problem, the other component of which (the choice of frame or reference point) cannot be observed.

Such models are hardly new to economic analysis. Indeed, most economic models presume that the agent is simultaneously maximizing with respect to variables other than the ones being studied. When assumptions are made on the individual's joint preferences over the observed and unobserved variables, the well-developed *theory of induced preferences*[29] can be used to derive testable implications on choice behavior over the observables. With a little more knowledge on exactly how frames are chosen, such an approach could presumably be applied here as well.

The above remarks should not be taken as implying that we have already solved the problem of framing in economic analysis or that there is no need to adapt, and if necessary abandon, our standard models in light of this phenomenon. Rather, they

---

[29] E.g. Milne (1981). For an application of the theory of induced preferences to choice under uncertainty, see Machina (1984).

reflect the view that when psychologists are able to hand us enough systematic evidence on how these effects operate, economists will be able to respond accordingly.

Other Issues: Is Probability Theory Relevant?

The Manipulation of Subjective Probabilities

The evidence discussed so far has primarily consisted of cases where subjects have been presented with explicit (i.e. "objective") probabilities as part of their decision problems, and the models which have addressed these phenomena possess the corresponding property of being defined over objective probability distributions. However, there is extensive evidence that when individuals have to estimate or revise probabilities for themselves they will make systematic mistakes in doing so.

The psychological literature on the processing of probabilistic information is much too large even to summarize here. However, it is worth noting that experimenters have uncovered several "heuristics" used by subjects which can lead to predictable errors in the formation and manipulation of subjective probabilities. Kahneman and Tversky (1973), Bar-Hillel (1974) and Grether (1980), for example, have found that probability updating systematically departs from Bayes Law in the direction of underweighting prior information and overweighting the "representativeness" of the current sample. In a related phenomenon termed the "law of small numbers," Tversky and Kahneman (1971) found that individuals overestimated the probability of drawing a perfectly representative sample out of a heterogeneous population. Finally, Bar-Hillel (1973), Tversky and Kahneman (1983) and others have found systematic biases in the formation of the probabilities of conjunctions of both independent and non-independent events. For surveys, discussions and examples of the psychological literature on the formation and handling of probabilities see Edwards, Lindman and Savage (1963), Slovic and Lichtenstein (1971), Tversky and Kahneman (1974) and the collections in *Acta Psychologica* (December 1970), Kahneman, Slovic and Tversky (1982) and Arkes and Hammond (1986). For examples of how economists have responded to some of these issues see Arrow (1982), Viscusi (1985) and the references cited there.

The Existence of Subjective Probabilities

The evidence referred to above indicates that when individuals are asked to formulate probabilities they do not do it correctly. However, these findings may be rendered moot by evidence which suggests that when individuals making decisions under uncertainty are not explicitly asked to form subjective probabilities, they might not do it (or even act as if doing it) at all.

In one of a class of examples due to Ellsberg (1961), subjects were presented with a pair of urns, the first containing 50 red balls and 50 black balls and the second also containing 100 red and black balls but in an unknown proportion. When faced with the choice of staking a prize on: (R_1) drawing a red ball from the first urn, (R_2)

drawing a red ball from the second urn, (B_1) drawing a black ball from the first urn, or (B_2) drawing a black ball from the second urn, a majority of subjects strictly preferred (R_1) over (R_2) and strictly preferred (B_1) over (B_2). It is clear that there can exist no subjectively assigned probabilities $p: (1 - p)$ of drawing a red vs. black ball from the second urn, even $1/2 : 1/2$, which can simultaneously generate both of these strict preferences. Similar behavior in this and related problems has been observed by Raiffa (1961), Becker and Brownson (1964), Slovic and Tversky (1974) and MacCrimmon and Larsson (1979).

Life (and Economic Analysis) Without Probabilities

One response to this type of phenomenon has been to suppose that individuals "slant" whatever subjective probabilities they might otherwise form in a manner which reflects the amount of confidence/ambiguity associated with them (Fellner, 1961; Becker and Brownson, 1964; Fishburn, 1986; Hogarth and Kunreuther, 1986). In the case of the complete ignorance regarding probabilities, Arrow and Hurwicz (1972), Maskin (1979) and others have presented axioms which imply principles such as ranking options solely on the basis of their worst and/or best outcomes (e.g. maximin, maximax), the unweighted average of their outcomes ("principle of insufficient reason"), or similar criteria.[30] Finally, generalizations of expected utility theory which drop the standard additivity and/or compounding laws of probability theory have been developed by Schmeidler (1986) and Segal (1987).

Although the above models may well capture aspects of actual decision processes, the analytically most useful approach to choice in the presence of uncertainty but the absence of probabilities is the so-called *state-preference* model of Arrow (1953/1964), Debreu (1959) and Hirshleifer (1966). In this model uncertainty is represented by a set of mutually exclusive and exhaustive *states of nature* $S = \{ s_i \}$. This partition of all possible unfoldings of the future could be either very coarse, such as the pair of states {it rains here tomorrow, it doesn't rain here tomorrow} or else very fine, so that the definition of a state might read "it rains here tomorrow and the temperature at Gibraltar is 75° at noon and the price of gold in New York is below $700.00/ounce." Note that it is neither feasible nor desirable to capture all conceivable sources of uncertainty when specifying the set of states for a given problem: it is not feasible since no matter how finely the states are defined there will always be some other random criterion on which to further divide them, and not desirable since such criteria may affect neither individuals' preferences nor their opportunities. Rather, the key requirements are that the states be mutually exclusive and exhaustive so that exactly one will be realized, and (for purposes of the present discussion) that the individual cannot influence which state will actually occur.

Given a fixed (and say finite) set of states, the objects of choice in this framework consist of alternative *state-payoff bundles*, each of which specifies the outcome the individual will receive in every possible state. When the outcomes are monetary payoffs, for example, state-payoff bundles take the form $(c_1, \ldots, c_n)$, where c_i denotes

---

[30] For an excellent discussion of the history, nature and limitations of such approaches, see Arrow (1951).

the payoff the individual would receive should state s_i occur. In the case of exactly two states of nature we could represent this set by the points in the (c_1, c_2) plane. Since bundles of the form (c, c) represent prospects which yield the same payoff in each state of nature, the $45°$ line in this plane is known as the *certainty line*.

Now if the individual did happen to assign probabilities $\{ p_i \}$ to the states $\{ s_i \}$, each bundle $(c_1, \ldots, c_n)$ would imply a specific probability distribution over wealth, and we could infer his or her preferences (i.e. indifference curves) over state-payoff bundles. However, since these bundles are defined directly over the respective states and without reference to any probabilities, it is also possible to speak of preferences over these bundles without making any assumptions regarding the coherency, or even existence, of such probabilistic beliefs. Researchers such as the ones cited above as well as Yaari (1969), Diamond and Yaari (1972) and Mishan (1976) have shown how this indifference curve-based approach can be used to derive results from individual demand behavior through general equilibrium in a context which requires neither the expected utility hypothesis nor the existence or commonality of subjective probabilities. In other words, life without probabilities does not imply life without economic analysis.

Final Thoughts

Welfare Implications.   Although the theme of this paper has been the descriptive theory of choice under uncertainty, another important issue is the implications of these developments for normative economics. Can welfare analysis be conducted in the type of world implied by the above models?

The answer to this question depends upon the model. Fanning-out behavior and the non-expected utility models used to characterize it, as well as the state-payoff approach of the previous section, are completely consistent with the assumption of well-defined, transitive individual preference orderings, and hence with traditional welfare analysis along the lines of Pareto, Bergson and Samuelson (e.g. Samuelson, 1947/1983, Ch. VIII). For example, the proof of Pareto-efficiency of a system of complete contingent-commodity markets (Arrow, 1953/1964; Debreu, 1959, Ch. 7) requires neither the expected utility hypothesis nor the assumption of well-defined probabilistic beliefs. On the other hand, it is clear that the preference reversal phenomenon and framing effects, and at least some of the non-transitive and/or non-economic models used to address them, will prove much more difficult to reconcile with welfare analysis, at least as currently practiced.

A Unified Model?   Another issue is the lack of a unified model capable of simultaneously handling all of the phenomena described in this paper: fanning-out, the preference reversal phenomenon, framing effects, probability biases and the Ellsberg paradox. After all, it is presumably the same ("typical") individuals who are exhibiting each of these phenomena—shouldn't there be a single model out there capable of generating them all?

Although I am doubtful of our present ability to do this, I am also doubtful about the need to establish a unified model as a prerequisite for continued progress. The aspects of behavior considered in this paper are very diverse, and if (like the wave

versus particle properties of light) they cannot be currently unified, this does not mean that we cannot continue to learn by studying and modelling them separately.

An Essential Criterion. The evidence and theories reported in this paper have taken us a long way from the classical expected utility approach presented at the outset. To what extent will these new models be incorporated into mainstream economic thought and practice? I believe the answer will depend upon a single factor: the extent to which they can address the important issues in the economics of uncertainty, such as search, investment, bargaining or auctions, to which the expected utility model has been so usefully applied.

References

Akerlof, George A., "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, August 1970, *84*, 488–500.

Allais, Maurice, "Le Comportement de l'Homme Rationel devant le Risque, Critique des Postulates et Axiomes de l'École Americaine," *Econometrica*, October 1953, *21*, 503–46.

Allais, Maurice, "The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School," in Allais and Hagen (1979).

Allais, Maurice and Ole Hagen, eds., *Expected Utility Hypotheses and the Allais Paradox*. Dordrecht, Holland: D. Reidel, 1979.

Allen, Beth, "Smooth Preferences and the Local Expected Utility Hypothesis," *Journal of Economic Theory*, 1987, forthcoming.

Arkes, Hal R. and Kenneth R. Hammond, eds., *Judgement and Decision Making: An Interdisciplinary Reader*. Cambridge: Cambridge University Press, 1986.

Arrow, Kenneth J., "Alternative Approaches to the Theory of Choice in Risk-Taking Situations," *Econometrica*, October 1951, *19*, 404–37. Reprinted in Arrow (1965).

Arrow, Kenneth J., "Le Role des Valeurs Boursières pour la Répartition le meilleure des risques," *Économetrie*, Colloques Internationaux du Centre National de la Recherche Scientifique, Paris, 1953, *40*, 41–47. English translation: *Review of Economic Studies*, April 1964, *31*, 91–96.

Arrow, Kenneth J., *Aspects of the Theory of Risk-Bearing*, Helsinki: Yrjo Jahnsson Saatio, 1965.

Arrow, Kenneth J., "Risk Perception in Psychology and Economics," *Economic Inquiry*, January 1982, *20*, 1–9.

Arrow, Kenneth J. and Leonid Hurwicz, "An Optimality Criterion for Decision-Making under Ignorance." In Carter, C. F. and J. L. Ford, eds., *Uncertainty and Expectations in Economics*. Oxford: Basil Blackwell, 1972.

Bar-Hillel, Maya, "On the Subjective Probability of Compound Events," *Organizational Behavior and Human Performance*, June 1973, *9*, 396–406.

Bar-Hillel, Maya, "Similarity and Probability," *Organizational Behavior and Human Performance*, April 1974, *11*, 277–82.

Battalio, Raymond C., John H. Kagel and Don N. MacDonald, "Animals' Choices over Uncertain Outcomes," *American Economic Review*, September 1985, *75*, 597–613.

Becker, Selwyn W. and Fred O. Brownson, "What Price Ambiguity? Or the Role of Ambiguity in Decision-Making," *Journal of Political Economy*, February 1964, *72*, 62–73.

Becker, Gordon M., Morris H. DeGroot and Jacob Marschak, "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, July 1964, *9*, 226–32.

Bell, David E., "Regret in Decision Making Under Uncertainty," *Operations Research*, September–October 1982, *30*, 961–81.

Bell, David E., "Disappointment in Decision Making Under Uncertainty," *Operations Research*, January–February 1985, *33*, 1–27.

Berg, Joyce E., John W. Dickhaut and John R. O'Brien, "Preference Reversal and Arbitrage," manuscript, University of Minnesota, September 1983.

Bernoulli, Daniel, "Specimen Theoriae Novae de Mensura Sortis," *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, 1738, *5*, 175–92. English translation: *Econometrica*, January 1954, *22*, 23–36.

Blyth, Colin R., "Some Probability Paradoxes in Choice from Among Random Alternatives," *Journal of the American Statistical Association*, June 1972, *67*, 366–73.

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

Chew Soo Hong, Edi Karni and Zvi Safra, "Risk Aversion in the Theory of Expected Utility with Rank Dependent Probabilities," *Journal of Economic Theory*, 1987, forthcoming.

Chew Soo Hong and William Waller, "Empirical Tests of Weighted Utility Theory," *Journal of Mathematical Psychology*, March 1986, *30*, 55–72.

Debreu, Gerard, *Theory of Value: An Axiomatic Analysis of General Equilibrium*. New Haven: Yale University Press, 1959.

Diamond, Peter A. and Menahem Yaari, "Implications of the Theory of Rationing for Consumer Choice Under Uncertainty," *American Economic Review*, June 1972, *62*, 333–43.

Edwards, Ward, "The Theory of Decision Making," *Psychological Bulletin*, July 1954, *51*, 380–417.

Edwards, Ward, "The Prediction of Decisions Among Bets," *Journal of Experimental Psychology*, September 1955, *50*, 201–14.

Edwards, Ward, Harold Lindman and Leonard J. Savage, "Bayesian Statistical Inference for Psychological Research," *Psychological Review*, May 1963, *70*, 193–242.

Ellsberg, Daniel, "Risk, Ambiguity and the Savage Axioms," *Quarterly Journal of Economics*, November 1961, *75*, 643–69.

Epstein, Larry, "Decreasing Risk Aversion and Mean-Variance Analysis," *Econometrica*, 1985, *53*, 945–61.

Fellner, William, "Distortion of Subjective Probabilities as a Reaction to Uncertainty," *Quarterly Journal of Economics*, November 1961, *75*, 670–89.

Fischhoff, Baruch, "Predicting Frames," *Journal of Experimental Psychology: Learning, Memory and Cognition*, January 1983, *9*, 103–16.

Fishburn, Peter C., "Nontransitive Measurable Utility," *Journal of Mathematical Psychology*, August 1982, *26*, 31–67.

Fishburn, Peter C., "Transitive Measurable Utility," *Journal of Economic Theory*, December 1983, *31*, 293–317.

Fishburn, Peter C., "SSB Utility Theory: An Economic Perspective," *Mathematical Social Sciences*, 1984, *8*, 63–94.

Fishburn, Peter C., "A New Model for Decisions Under Uncertainty," *Economics Letters*, 1986, *21*, 127–30.

Friedman, Milton and Leonard J. Savage, "The Utility Analysis of Choices Involving Risk," *Journal of Political Economy*, August 1948, *56*, 279–304.

Grether, David M., "Bayes Rule as a Descriptive Model: The Representativeness Heuristic," *Quarterly Journal of Economics*, November 1980, *95*, 537–57.

Grether, David M. and Charles R. Plott, "Economic Theory of Choice and the Preference Reversal Phenomenon," *American Economic Review*, September 1979, *69*, 623–38.

Hagen, Ole, "Towards a Positive Theory of Preferences Under Risk," in Allais and Hagen (1979).

Hershey, John C. and Paul J. H. Schoemaker, "Risk-Taking and Problem Context in the Domain of Losses—An Expected Utility Analysis," *Journal of Risk and Insurance*, March 1980 *47*, 111–32.

Hershey, John C. and Paul J. H. Schoemaker, "Probability Versus Certainty Equivalence Methods in Utility Measurement: Are They Equivalent?," *Management Science*, October 1985, *31*, 1213–31.

Hirshleifer, Jack, "Investment Decision Under Uncertainty: Applications of the State-Preference Approach," *Quarterly Journal of Economics*, May 1966, *80*, 252–77.

Hogarth, Robin, "Cognitive Processes and the Assessment of Subjective Probability Distributions," *Journal of the American Statistical Association*, June 1975, *70*, 271–89.

Hogarth, Robin and Howard Kunreuther, "Decision Making Under Ambiguity," *Journal of Business*, October 1986, Prt. 2, *4*, 225–50.

Holt, Charles A., "Preference Reversals and the Independence Axiom," *American Economic Review*, June 1986, *76*, 508–15.

Kahneman, Daniel, Paul Slovic and Amos Tversky, eds., *Judgement Under Uncertainty: Heuristics and Biases*, Cambridge: Cambridge University Press, 1982.

Kahneman, Daniel and Amos Tversky, "On the Psychology of Prediction," *Psychological Review*, July 1973, *80* 237–51.

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

Karmarkar, Uday S., "The Effect of Probabilities on the Subjective Evaluation of Lotteries," Massachusetts Institute of Technology Sloan School of Business Working Paper, 1974.

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

Karni, Edi and Zvi Safra, "'Preference Reversal' and the Observability of Preferences by Experimental Methods," *Econometrica*, 1987, forthcoming.

Keller, L. Robin, "The Effects of Decision Problem Representation on Utility Conformance," manuscript, University of California, Irvine, 1982.

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

Knez, Marc and Vernon L. Smith, "Hypothetical Valuations and Preference Reversals in the Context of Asset Trading," manuscript, University of Arizona, 1986.

Lichtenstein, Sarah and Paul Slovic, "Reversals of Preferences Between Bids and Choices in Gambling Decisions," *Journal of Experimental Psychology*, July 1971, *89*, 46–55.

Lichtenstein, Sarah and Paul Slovic, "Response-Induced Reversals of Preference in Gambling: An Extended Replication in Las Vegas," *Journal of Experimental Psychology*, November 1973, *101*, 16–20.

Lindman, Harold, "Inconsistent Preferences Among Gambles," *Journal of Experimental Psychology*, May 1971, *89*, 390–97.

Loomes, Graham and Robert Sugden, "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty," *Economic Journal*, December 1982, *92*, 805–24.

MacCrimmon, Kenneth R., "Descriptive and Normative Implications of the Decision-Theory Postulates." In Borch Karl H., and Jan Mossin, eds., *Risk and Uncertainty: Proceedings of a Conference Held by the International Economic Association*. London: Macmillan, 1968.

MacCrimmon, Kenneth R. and Stig Larsson, "Utility Theory: Axioms Versus 'Paradoxes,'" in

Allais and Hagen (1979).

MacCrimmon, Kenneth R. and Donald A. Wehrung, *Taking Risks: The Management of Uncertainty*, New York: The Free Press, 1986.

Machina, Mark J., "'Expected Utility' Analysis Without the Independence Axiom," *Econometrica*, March 1982, *50*, 277–323.

Machina, Mark J., "The Economic Theory of Individual Behavior Toward Risk: Theory, Evidence and New Directions," Stanford University Institute for Mathematical Studies in the Social Sciences Technical Report, 1983a.

Machina, Mark J., "Generalized Expected Utility Analysis and the Nature of Observed Violations of the Independence Axiom," 1983b, in Stigum and Wenstøp (1983).

Machina, Mark J., "Temporal Risk and the Nature of Induced Preferences," *Journal of Economic Theory*, August 1984, *33*, 199–231.

Markowitz, Harry "The Utility of Wealth," *Journal of Political Economy*, April 1952, *60*, 151–58.

Marschak, Jacob, "Rational Behavior, Uncertain Prospects, and Measurable Utility," *Econometrica*, April 1950, *18*, 111–41. "Errata," *Econometrica*, July 1950, *18*, 312.

Mas-Colell, Andreu, "An Equilibrium Existence Theorem Without Complete or Transitive Preferences," *Journal of Mathematical Economics*, December 1974, *3*, 237–46.

Maskin, Eric, "Decision Making Under Ignorance with Implications for Social Choice," *Theory and Decision*, September 1979, *11*, 319–37.

May, Kenneth O., "Intransitivity, Utility, and the Aggregation of Preference Patterns," *Econometrica*, January 1954, *22*, 1–13.

McCord, Marc and Richard de Neufville, "Empirical Demonstration that Expected Utility Analysis Is Not Operational." In Stigum and Wenstøp (1983).

McCord, Marc and Richard de Neufville, "Utility Dependence on Probability: An Empirical Demonstration," *Large Scale Systems*, February 1984, *6*, 91–103.

McNeil, Barbara J., Stephen G. Pauker, Harold C. Sox, Jr. and Amos Tversky, "On the Elicitation of Preferences for Alternative Therapies," *New England Journal of Medicine*, May 1982, *306*, 1259–62.

Milne, Frank, "Induced Preferences and the Theory of the Consumer," *Journal of Economic Theory*, April 1981, *24*, 205–17.

Mishan, E. J., "Choices Involving Risk: Simple Steps Toward an Ordinalist Analysis," *Economic Journal*, December 1976, *86*, 759–77.

Morrison, Donald G., "On the Consistency of Preferences in Allais' Paradox," *Behavioral Science*, September 1967, *12*, 373–83.

Moskowitz, Herbert, "Effects of Problem Rep-

resentation and Feedback on Rational Behavior in Allais and Morlat-Type Problems," *Decision Sciences*, 1974, *5*, 225–42.

Mowen, John C. and James W. Gentry, "Investigation of the Preference-Reversal Phenomenon in a New Product Introduction Task," *Journal of Applied Psychology*, December 1980, *65*, 715–22.

Payne, John W. and Myron L. Braunstein, "Preferences Among Gambles with Equal Underlying Distributions," *Journal of Experimental Psychology*, January 1971, *87*, 13–18.

Pommerehne, Werner W., Friedrich Schneider and Peter Zweifel, "Economic Theory of Choice and the Preference Reversal Phenomenon: A Reexaminaton," *American Economic Review*, June 1982, *72*, 569–74.

Pratt, John W., "Risk Aversion in the Small and in the Large," *Econometrica*, January/April 1964, *32*, 122–36.

Quiggin, John, "A Theory of Anticipated Utility," *Journal of Economic Behavior and Organization*, December 1982, *3*, 323–43.

Raiffa, Howard, "Risk, Ambiguity, and the Savage Axioms," *Quarterly Journal of Economics*, November 1961, *75*, 690–94.

Raiffa, Howard, *Decision Analysis: Introductory Lectures on Choice Under Uncertainty*. Reading, Mass.: Addison-Wesley, 1968.

Reilly, Robert J., "Preference Reversal: Further Evidence and Some Suggested Modifications of Experimental Design," *American Economic Review*, June 1982, *72*, 576–84.

Rothschild, Michael and Joseph E. Stiglitz, "Increasing Risk: I. A Definition," *Journal of Economic Theory*, September 1970, *2*, 225–43.

Russo, J. Edward, "The Value of Unit Price Information," *Journal of Marketing Research*, May 1977, *14*, 193–201.

Russo, J. Edward, Gene Krieser and Sally Miyashita, "An Effective Display of Unit Price Information," *Journal of Marketing*, April 1975, *39*, 11–19.

Samuelson, Paul A., *Foundations of Economic Analysis*, Cambridge, Mass.: Harvard University Press, 1947. Enlarged Edition, 1983.

Samuelson, Paul A., *Foundations of Economic Analysis*, Cambridge, Mass.: Harvard University Press, 1947. Enlarged Edition, 1983.

Savage, Leonard J., *The Foundations of Statistics*, New York: Wiley, 1954. Revised and Enlarged Edition, New York: Dover, 1972.

Schmeidler, David, "Subjective Probability and Expected Utility Without Additivity," manuscript, Tel-Aviv University, 1986.

Schoemaker, Paul J. H. and Howard Kunreuther, "An Experimental Study of Insurance Decisions," *Journal of Risk and Insurance*, December 1979, *46*, 603–18.

Segal, Uzi, "Nonlinear Decision Weights with the Independence Axiom," manuscript, University of California, Los Angeles, November 1984.

Segal, Uzi, "The Ellsberg Paradox and Risk Aversion: An Anticipated Utility Approach," *International Economic Review*, 1987, forthcoming.

Shafer, Wayne J., "The Nontransitive Consumer," *Econometrica*, September 1974, *42*, 913–19.

Slovic, Paul, "Manipulating the Attractiveness of a Gamble Without Changing its Expected Value," *Journal of Experimental Psychology*, January 1969, *79*, 139–45.

Slovic, Paul, "Choice Between Equally Valued Alternatives," *Journal of Experimental Psychology: Human Perception and Performance*, August 1975, *1*, 280–87.

Slovic, Paul, Baruch Fischhoff and Sarah Lichtenstein, "Response Mode, Framing, and Information Processing Effects in Risk Assessment." In Hogarth, Robin ed., *New Directions for Methodology of Social and Behavioral Science: Question Framing and Response Consistency*. San Francisco: Jossey-Bass, 1982.

Slovic, Paul, Baruch Fischoff and Sarah Lichtenstein, "Relative Importance of Probabilities and Payoffs in Risk Taking," *Journal of Experimental Psychology*, November 1968, Prt. 2, *78*, 1–18.

Slovic, Paul, Baruch Fischoff and Sarah Lichtenstein, "Comparison of Bayesian and Regression Approaches to the Study of Information Processing in Judgment," *Organizational Behavior and Human Performance*, November 1971, *6*, 649–744.

Slovic, Paul, Baruch Fischoff and Sarah Lichtenstein, "Preference Reversals: A Broader Perspective," *American Economic Review*, September 1983, *73*, 596–605.

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

Sonnenschein, Hugo F., "Demand Theory Without Transitive Preferences, With Applications to the Theory of Competitive Equilibrium," In Chipman, John S., Leonid Hurwicz, Marcel K. Richter, and Hugo F. Sonnenschein, eds., *Preferences, Utility and Demand*. New York: Harcourt Brace Jovanovich, 1971.

Spence, A. Michael and Richard J. Zeckhauser, "Insurance, Information and Individual Action," *American Economic Review Papers and Proceedings*, May 1971, *61*, 380–87.

Stigum, Bernt and Fred Wenstøp, *Foundations of Utility and Risk Theory with Applications*, Dordrecht, Holland: D. Reidel, 1983.

Sugden, Robert, "New Developments in the Theory of Choice Under Uncertainty," *Bulletin of*

Economic Research, January 1986, *38*, 1–24.

Thaler, Richard, "Toward a Positive Theory of Consumer Choice," *Journal of Economic Behavior and Organization*, March 1980, *1*, 39–60.

Tversky, Amos, "Intransitivity of Preferences," *Psychological Review*, January 1969, *76*, 31–48.

Tversky, Amos, "A Critique of Expected Utility Theory: Descriptive and Normative Considerations," *Erkenntnis*, 1975, *9*, 163–73.

Tversky, Amos and Daniel Kahneman, "Belief in the Law of Small Numbers," *Psychological Bulletin*, July 1971, *2*, 105–10.

Tversky, Amos and Daniel Kahneman, "Judgement under Uncertainty: Heuristics and Biases," *Science*, September 1974, *185*, 1124–31.

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

Tversky, Amos and Daniel Kahneman, "Ex-tensional vs. Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment," *Psychological Review*, October 1983, *90*, 293–315.

Tversky, Amos and Daniel Kahneman, "Rational Choice and the Framing of Decisions," *Journal of Business*, October 1986, Prt. 2, *4*, 251–78.

Viscusi, W. Kip, "Are Individuals Bayesian Decision Makers?" *American Economic Review Papers and Proceedings*, May 1985, *75*, 381–85.

von Neumann, John and Oskar Morgenstern, *Theory of Games and Economic Behavior*. Princeton: Princeton University Press, 1944. 2nd Ed., 1947. 3rd Ed., 1953.

Yaari, Menahem, "Some Remarks on Measures of Risk Aversion and On their Uses," *Journal of Economic Theory*, October 1969, *1*, 315–29.

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