WILLIAM S. NEILSON

PROBABILITY TRANSFORMATIONS IN THE STUDY OF BEHAVIOR TOWARD RISK

ABSTRACT. Probability transformation functions were introduced into models of behavior toward risk to allow them to accommodate violations of the expected utility hypothesis. This paper examines the shape of the probability transformation function, its interpretation as optimism or pessimism, and how the ranking of outcomes becomes important when probability transformations are used. It also explores two behavioral implications: the overweighting of unlikely, extreme outcomes, and intertia around certainty. Finally, the rationality of transforming the probability distribution is discussed.

1. INTRODUCTION

In its original form, prospect theory (Kahneman and Tversky 1979) makes two important changes to the expected utility model of rational choice. One is the idea that individuals care not about final wealth levels, but rather about changes from some reference level of wealth, with losses treated differently from gains. The other papers in this symposium deal with the implications of this idea. The second contribution of prospect theory is that for the purposes of representing preferences, the mathematical expectation of the utility values of the outcomes should not be based on the true, objective probabilities, but instead that the probabilities should be transformed before the expectation is computed. The nature of these probability transformations, the evidence pertaining to them, and their behavioral implications are the focus of this paper.

Oddly enough, the different social sciences have focused on different aspects of prospect theory. Research within political science and sociology has concentrated almost exclusively on the reference dependence of utility. In contrast, most of the research related to prospect theory within the economics literature concerns issues pertaining to probability transformations. Only psychologists have paid considerable attention to both issues. This fact that most of the research on probability transformations has been performed by economists and psychologists means that little will be said about political science or sociology in this paper.

Synthese **135:** 171–192, 2003.

The intuition behind probability transformations is quite simple. Experimental evidence makes it apparent that individuals tend to overweight unlikely, extreme outcomes and underweight likely, extreme outcomes, relative to their objective probabilities. Expected utility, by definition, weights outcomes by their objective probabilities, so it is inconsistent with this experimental evidence. Preference representations that employ probability transformations are simply an effort to accommodate these violations of expected utility.

Somewhat ironically, the most frequently-used preference representation that employs probability transformations is not the one proposed by Kahneman and Tversky in their original prospect theory paper. Instead, the state-of-the-art stems from the work of Quiggin (1982), Yaari (1987), and Schmeidler (1989). The original prospect theory formulation called for each probability to be transformed individually, whereas the other representations call for the entire probability distribution to be transformed at once. The resulting model is known as rank-dependent utility. In their revision of prospect theory, Tversky and Kahneman (1992) utilize rank-dependent utility along with reference dependence.

It is worth stating at the outset what proper and improper interpretations of probability transformations are. The goal of a preference representation is to construct a mathematical formula that is capable of either predicting or accommodating the choices of an individual in a specific setting. With the rank-dependent utility representation, the mathematical formula states that the (monetary) outcomes are transformed by a utility function, the probability distribution is transformed by a transformation function, and the mathematical expectation of the utility values is taken using the transformed probabilities. This formulation does not imply that individuals incorrectly assess monetary payoffs; rather, it simply uses the utility function as a tool for modeling preferences. Likewise, the formulation does not imply that individuals incorrectly assess probabilities; rather, it simply uses the transformation function as a tool for representing preferences.

The paper proceeds as follows. Section 2 introduces rank-dependent utility, compares it to expected utility, and shows that it meets the standard minimal rationality requirement for preferences toward risk – an ability to allow a preference for first-order stochastically dominating shifts in probability distributions. Section 3 goes on to discuss properties of the probability transformation function, and introduces new notions of risk attitudes – optimism and pessimism. Sections 4 and 5 discuss behavioral implications. Section 4 covers implications of the shape of the probability transformation function, and these implications arise from the overweighting of extreme, unlikely events. Section 5 covers implications of rank-dependence, which can be interpreted as inertia around certainty. The paper concludes in Section 6 with a discussion of the notion of rationality, as it pertains to preferences toward risk, in light of the results of the rank-dependent utility model.

2. THE FORM OF RANK-DEPENDENT UTILITY

Let X be a finite subset of the real numbers, $\{x_1, \dots, x_n\}$, and let $P(X)$ be the set of probability distributions defined over X. Each $x \in X$ is a *prize* or *outcome*, and each distribution $p \in P(X)$ is a *lottery*. The lottery p is sometimes written $p = (p_1, \dots, p_n)$, where p_i is the probability of the outcome x_i. An individual has preferences defined over $P(X)$. The basic task of utility theory is to construct a function that represents preferences; that is, to construct a *preference function* $V\colon P \to R$ such that $V(p) \geq V(q)$ if and only if the lottery p is weakly preferred to the lottery q. Typically these preference functions are held to either normative or positive standards, or both. For example, expected utility has normative appeal because of the axioms it satisfies, and it has positive appeal because it is able to describe behavior in a wide variety of settings.

In this section I introduce the preference function associated with rank-dependent utility. Before doing so, it is useful to characterize expected utility, both as a point of departure and as a point of comparison. When the lottery has n outcomes, expected utility takes the form

$$EU(p) = \textstyle\sum_{i=1}^{n} p_i u(x_i), \tag{1}$$

where $u(x)$ is the *utility function*. The existence of a utility function has proven to be extremely useful for several reasons. First, as shown by Arrow (1974) and Pratt (1964), for example, risk attitudes can be captured by the curvature of the utility function, with a concave utility function corresponding to risk aversion, a convex utility function corresponding to risk seeking, and a linear utility function corresponding to risk neutrality. Second, and perhaps more importantly, the utility function allows the researcher to govern attitudes toward risk in an $(n-1)$-dimensional space (i.e., the probability space $P(X)$) using a one-dimensional function. This property becomes especially important when one is concerned with continuous probability distributions instead of discrete ones, because then the utility function allows the researcher to collapse an infinite-dimensional space into a single dimension.

The expected utility preference function in (1) can be rewritten in a way that clarifies the presentation of rank-dependent utility. Assume, without loss of generality, that the members of X are numbered so that $x_1 < \cdots < x_n$. Rewriting (1) yields

$$EU(p) = p_1 u(x_1) + \textstyle\sum_{i=2}^{n}[(p_1 + \cdots + p_i) - (p_1 + \cdots + p_{i-1})]u(x_i). \tag{2}$$

Letting F denote the standard cumulative distribution function, so that $F(x_0) = \text{Prob}\{x \leq x_0\}$, Equation (2) can be further rewritten as

$$EU(p) = F(x_1)u(x_1) + \textstyle\sum_{i=2}^{n}[F(x_i) - F(x_{i-1})]u(x_i). \tag{3}$$

As Equation (1) shows, in expected utility the "decision weight" is simply the probability of the prize, or, as Equation (3) shows, the decision weight is also the difference between the cumulative probability of the prize, $F(x_i)$, and the cumulative probability of the next-best prize, $F(x_{i-1})$.

Rank-dependent utility differs from expected utility by utilizing a different scheme of decision weights. The rank-dependent utility preference function is given by

$$V(p) = h(p_1)u(x_1) + \textstyle\sum_{i=2}^{n}[h(p_1 + \cdots + p_i) - h(p_1 + \cdots + p_{i-1})]u(x_i), \tag{4}$$

or, using the cumulative distribution function notation,

$$V(p) = h(F(x_1))u(x_1) + \textstyle\sum_{i=2}^{n}[h(F(x_i)) - h(F(x_{i-1}))]u(x_i), \tag{5}$$

where h: [0, 1] $\to$ [0, 1] is a *probability transformation function* with the properties that it is strictly increasing with $h(0) = 0$ and $h(1) = 1$.

The only difference between expected utility and rank-dependent utility is that rank-dependent utility transforms the probability distribution before computing the expectation of utility. Where expected utility uses the (discrete) probability distribution F, rank-dependent utility uses the transformed probability distribution $h(F)$. Since h is increasing with $h(0) = 0$ and $h(1) = 1$, $h(F)$ is also a probability distribution, and so expectations are well-defined.

To see how the probability transformation function makes a difference, consider the special case of lotteries with just two outcomes, $x_1 < x_2$. The rank-dependent utility of the lottery p is

$$V(p) = h(p_1)u(x_1) + [1 - h(p_1)]u(x_2). \tag{6}$$

If $h(p_1) > p_1$, the individual overweights the lower outcome, x_1, relative to expected utility, making the lottery less attractive. Similarly, if

$h(p_1) < p_1$, the individual underweights the lower outcome relative to expected utility, and therefore overweights the larger outcome, making the lottery more attractive. In general, and letting $E[\cdot]$ denote the expectation operator, if $E[h(p)] > E[p]$, the individual is said to be *optimistic about the lottery* p, and if $E[h(p)] \geq E[p]$ for all lotteries in P with the strict inequality holding for at least one p, the individual is said to be *optimistic*. Similarly, if $E[h(p)] < E[p]$, the individual is said to be *pessimistic about the lottery* p, and if $E[h(p)] \leq E[p]$ for all lotteries in P with the strict inequality holding for at least one p, the individual is said to be *pessimistic*. Optimism and pessimism can work either with or against the risk aversion embodied by the utility function u, leading to different behavioral patterns than expected utility does. Some of these are investigated in the next section.

Before going on to discuss specific properties of rank-dependent utility and its probability transformation function, it is enlightening to investigate why the function transforms the probability distribution instead of the individual probabilities. Consider for the moment the preference function

$$W(p) = \sum_{i=1}^{n} g(p_i)u(x_i), \tag{7}$$

where g: $[0, 1] \to [0, 1]$ is an increasing, onto function. Such a functional form has been proposed by Handa (1977), Karmarker (1978), and Kahneman and Tversky (1979), although it is now out of fashion. In fact, even though prospect theory was originally formulated using a functional form like that in (7), Tversky and Kahneman (1992) revised prospect theory to give it a rank-dependent utility form, and named the new version *cumulative prospect theory*. The reason is that the functional form in (7) implies that individuals must dislike some first-order stochastically dominating shifts unless g is linear. A *first-order stochastically dominating shift* is a rightward shift of probability mass, and so individuals with expected utility preferences and increasing utility functions should prefer them.

To see the problem with (7), suppose that there are just two outcomes, $x_1 < x_2$, let $0 < \alpha < \beta < 1$, and construct four probability distributions, p, q, p' and q', with $p_1 = \alpha$, $q_1 = \beta$, $p'_2 = \beta$, and $q'_2 = \alpha$. Then p first-order stochastically dominates q since it places less weight on the low outcome than q does, and p' first-order stochastically dominates q' since it places more weight on the high outcome than q does. The individual prefers p to q if and only if

$$g(\alpha)u(x_1) + g(1 - \alpha)u(x_2) > g(\beta)u(x_1) + g(1 - \beta)u(x_2), \tag{8}$$

which can be rearranged to yield

$$\frac{g(\beta) - g(\alpha)}{g(1 - \alpha) - g(1 - \beta)} < \frac{u(x_2)}{u(x_1)}. \tag{9}$$

Similarly, the individual prefers p' to q' if and only if

$$(10) \qquad \frac{g(1-\alpha)-g(1-\beta)}{g(\beta)-g(\alpha)} < \frac{u(x_2)}{u(x_1)}.$$

As x_1 approaches x_2, the right-hand sides of both (9) and (10) approach one from above. Consequently, the left-hand sides of both expressions can be no greater than one. Since they are inverses of each other, they must be equal to one. The final step is to notice that $(1-\alpha)-(1-\beta) = \beta-\alpha$. So, the left-hand side of (10) being equal to one implies that

$$(11) \qquad \frac{g(\beta)-g(\alpha)}{\beta-\alpha} = \frac{g(1-\alpha)-g(1-\beta)}{(1-\alpha)-(1-\beta)},$$

which in turn implies that g has a constant slope, which can only occur if $g(p) = p$.

Rank-dependent utility avoids this problem because the probability transformation is applied to the distribution function, not the individual probabilities. In fact, if both u and g are increasing, the rank-dependent utility preference function must exhibit a preference for first-order stochastically dominating shifts. Using the same lotteries p and q as above, with p first-order stochastically dominating q, yields

$$(12) \qquad V(p) = h(\alpha)u(x_1) + [1-h(\alpha)]u(x_2),$$

and

$$(13) \qquad V(q) = h(\beta)u(x_1) + [1-h(\beta)]u(x_2).$$

Since $\alpha < \beta$ and h is increasing, $h(\alpha) \leq h(\beta)$ and therefore $V(p)$ places greater weight on the higher utility value, $u(x_2)$, then $V(q)$ does. Consequently, $V(p)$ must be larger.[1]

The ability of a preference representation to accommodate a preference for first-order stochastically dominating shifts is a minimal rationality requirement. In a nonstochastic setting, this would correspond to the ability of a preference representation to accommodate a preference for increases in the size of the consumption bundle. Note that the desired property is a representation that *can* handle a preference for increases in consumption, not a representation that *must* imply a preference for increases in consumption. There are situations in which an individual might dislike increases in consumption, due to satiation, for example. But, if the representation is unable to accommodate a preference for increases in consumption, then it must be the case that there are increases in consumption that the individual *must* dislike. This is a very strong restriction on preferences, and one that researchers are unwilling to consider rational. The same arguments hold for preferences for first-order stochastically dominating shifts. If a representation implies that the decision-maker must dislike some first-order stochastically dominating shifts, that representation is usually discarded on rationality grounds. Rank-dependent utility fits this minimal rationality requirement.

3. THE PROBABILITY TRANSFORMATION FUNCTION

In Section 2, rank-dependent utility was presented as a mathematical construct. This is not how it originated, however. Rank-dependent utility was created as the solution to the puzzle that arose from experimental violations of expected utility. The best-known violation is the Allais paradox. In the Allais paradox, individuals make choices in two pairs of hypothetical lotteries. Lottery A gives the individual \$1 million with certainty, while lottery B gives the individual \$5 million with probability 0.10, \$1 million with probability 0.89, and \$0 with probability 0.01. Lottery C pays the individual \$1 million with probability 0.11 and pays \$0 otherwise, and lottery D pays \$5 million with probability 0.10 and pays \$0 otherwise. A majority of experimental subjects prefer A to B, a majority prefer D to C, and the modal choice pair is A and D. The puzzle arises because this modal preference violates expected utility.

To see why, suppose that u is any utility function. An expected utility maximizer prefers lottery A to lottery B if

$$u(1\text{M}) > 0.01u(0) + 0.89u(1\text{M}) + 0.10u(5\text{M}), \tag{14}$$

which reduces to

$$0.11u(1\text{M}) > 0.01u(0) + 0.10u(5\text{M}). \tag{15}$$

Lottery D is preferred to lottery C if

$$0.90u(0) + 0.10u(5\text{M}) > 0.89u(0) + 0.11u(1\text{M}), \tag{16}$$

which can be rewritten

$$0.01u(0) + 0.10u(5\text{M}) > 0.11u(1\text{M}). \tag{17}$$

Obviously, there is no utility function for which both (15) and (17) can be true, so expected utility is violated.

The standard explanation for making the expected-utility violating choices in the Allais paradox is that, when deciding between lotteries A and B, the individual is unwilling to forego a sure \$1 million and risk getting \$0. Put another way, even though it also has a 10% chance of a much higher outcome (\$5 million), the individual is unwilling to choose B because the probability of getting zero is too large. When deciding between lotteries C and D, the individual views a positive outcome as unlikely in both cases, but the increase in the size of the payoff in D outweighs the small decrease in the probability. In other words, the 0.01 increase in the probability of receiving zero is outweighed by the increase in the payoff if the individual wins. The 0.01 increase in the probability of getting \$0 figures into both decisions, but it is weighed more heavily in the choice between A and B than in the choice between C and D.

Rank-dependent utility handles this asymmetric treatment of the 0.01 probability increase very easily. With rank-dependent utility, lottery A is preferred to lottery B if

$$u(1M) > h(0.01)u(0) + [h(0.90 - h(0.01)]u(1M) + [1 - h(0.90)]u(5M), \tag{18}$$

and the decision weight placed on the 0.01 increase in the probability of \$0 is $h(0.01)$. Lottery D is preferred to lottery C if

$$h(0.90)u(0) + [1 - h(0.90)]u(5M) > h(0.89)u(0) + [1 - h(0.89)]u(1M). \tag{19}$$

Here the decision weight placed on the 0.01 increase in the probability of \$0 is $h(0.90) - h(0.89)$. If $h(0.01) > h(0.90) - h(0.89)$, the Allais paradox can arise.[2]

The Allais paradox choices seem to suggest that $h(p) \geq p$, at least for small values of p. Researchers, mostly psychologists, have undertaken a more systematic approach to determine the shape of the probability transformation function. There are two important issues relevant to using experimental data to determine the probability transformation function, and I briefly discuss them here.

First, any attempt to uncover the shape of the probability transformation function using experimental evidence must keep the task performed by the subjects fixed. Psychologists have long argued that changing the task biases the results (Hershey et al. 1992). Perhaps the best example of this is the preference reversal phenomenon (Lichtenstein and Slovic 1971; Grether and Plott 1979; Tversky et al. 1990). Subjects are faced with a P-bet that assigns a high probability to a moderate outcome, and a \$-bet that assigns a low probability to a high outcome.[3] In one task each subject states a preference between the two gambles. In the other task each subject assigns a value to each of the two gambles. All utility theories predict that the preferred gamble should have a higher value. A large body of experimental evidence shows that this prediction fails. A large number of subjects prefer the P-bet to the $-bet but assign a higher value to the $-bet. Although these findings raise some serious issues about utility theories, they are not considered here. Instead, the preference reversal evidence is taken as a warning that experimental evidence must be used very carefully. Evidence can only be aggregated when the tasks are the same, and the results should only be interpreted within the context of the task.[4]

Second, there is some disagreement among the different social sciences about whether subjects should be paid on the basis of their choices or not. Economists, who tend to be more cognizant of incentives, argue forcefully that they should, and it is unusual for an experiment that does not pay subjects based on choices to be published in economics journals. They argue that making good choices requires some mental effort, and without payments based on choices, subjects will not use much mental effort and the results will be unrepresentative of behavior in situations where choices matter. Researchers in the other social sciences remain unconvinced. Some economists have performed experiments where some subjects were paid on the basis of their choices and others were paid a flat fee (Camerer 1989; Conlisk 1989; Battalio et al. 1990). The results tend to show that the direction of preferences seem to be the same with both real and hypothetical payoffs, but that the strength of preference and the standard errors are different. For example, Battalio et al. find that, quantitatively, subjects responding to real payoffs tend to be slightly more risk averse than subjects responding to hypothetical payoffs, but that qualitative conclusions based on the two settings tend to be the same.

With a few exceptions, the experimental data used to determine the shape of the probability transformation function is based on a single task, pairwise choices, and was generated using real payoffs. Typically, subjects are given a large number of choice pairs and instructed to state their preference in each pair. One of the choice pairs is selected at random, the two gambles are played, and each subject is paid according to the outcome of the gamble he or she chose. The gambles usually involve amounts of money in the $10 to $30 range for the highest payoff, large enough so that subjects take all of the choices seriously even though only one of them will matter.

The first efforts to attempt to identify the shape of the probability transformation function used a single parameterized form of the function (see

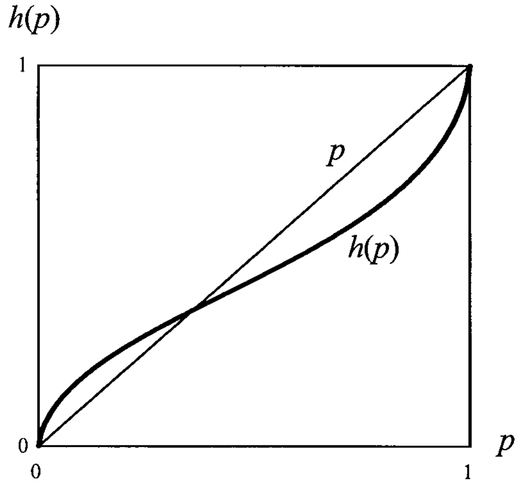

Figure 1.

Tversky and Kahneman 1992; Camerer and Ho 1994; Wu and Gonzalez 1996):

$$h(p) = \frac{p^{\gamma}}{(p^{\gamma} + (1 - p^{\gamma})^{\frac{1}{\gamma}}}. \tag{20}$$

For values of γ between zero and one, this function takes the form shown in Figure 1. Notice that when p is sufficiently small, h overweights the probability, capturing pessimism, and when p is large, h underweights the probability, which engenders optimism. One problem with this approach is that the functional form itself rules out certain types of behavior. For example, it is impossible for this functional form to have $h(p) \geq p$ for all values of p. So, a non-parametric approach might be better.

A non-parametric approach is taken by Gonzalez and Wu (1999), Bleichrodt and Pinto (2000), and Abdellaoui (2000). All three studies find that the general pattern of Figure 1 is robust. Subjects tend to overweight small probabilities and underweight large ones, with the crossover point lying somewhere between $p = 0.25$ and $p = 0.5$.

A completely different approach to the problem is taken by Prelec (1998). He devises a set of axioms which imply that the probability transformation function takes a specific functional form:

$$h(p) = \exp(-(-\ln p)^{\alpha}), \tag{21}$$

which has the same general shape as that shown in Figure 1. In Prelec's model, no matter what the value of α, the crossover point must be at $p =$

$1/e \approx 0.37$. So, according to the probability transformation function in (21), probabilities less than 0.37 are overweighted and probabilities greater than 0.37 are underweighted.

It is compelling that several experimental datasets, two different estimation approaches, and an approach based on normative theoretical axioms all lead to the same general patterns. Because of this, in the remainder of the paper the probability transformation function will be assumed to have the general shape shown in Figure 1. In terms of optimism and pessimism, when the lottery under consideration has only two outcomes, the individual is pessimistic about a lottery when the probability of the low outcome is below the crossover point, and he is optimistic about a lottery when the probability of the low outcome is above the crossover point. He can be neither everywhere pessimistic nor everywhere optimistic.

This is a rather strange pattern of optimism and pessimism. People are optimistic when success is unlikely, and pessimistic when failure is unlikely, which seems backwards. Is this a property that is genetically burned into people's decision-making apparatus, or is it learned? Insight into this question is provided by an experiment on elementary school children, performed by Harbaugh et al. (2002). They find that children use a probability transformation function, but one with a pattern opposite that as Figure 1. Children tend to underweight low probabilities and overweight high ones, which means that children are optimistic when success is likely and pessimistic when failure is likely. The Figure 1 pattern found for adults must be learned, then. This raises another interesting question – is the pattern in Figure 1 cultural or universal? If the pattern is learned, people in different cultures might exhibit different patterns, and there may be something particularly western about overweighting small probabilities, being optimistic when failure is likely, and being pessimistic when success is likely. As yet, there is no research on this topic.

4. IMPLICATIONS I – OVERWEIGHTING UNLIKELY, EXTREME OUTCOMES

The probability transformation function depicted in Figure 1 has an important property that affects predicted behavior: it overweights unlikely, extreme outcomes. In the case of an extreme loss, the decision weight placed on the utility of the loss is the transformed probability of the loss, which is higher than the true probability. In the case of an extreme gain, the decision weight placed on the utility of the gain is one minus the transformed cumulative probability of all lower outcomes, and since this cumulative probability is large, it is underweighted. Consequently, the extreme gain is overweighted.

The most straightforward application of this pattern is to insurance and gambling. It is readily apparent that people insure against large, unlikely losses but not against likely ones. For example, most people would be willing to pay $1 to insure against a 1/10,000 chance of losing $10,000, but very few people would be willing to pay $9000 to insure against a 9/10 chance of losing $10,000. Rank-dependent utility can help explain this. People buy insurance when the reduction in risk is worth the price of the premium. When the probability of the loss is small, the loss is overweighted, making the purchase of insurance more attractive. In contrast, when the probability of the loss is large, the loss is underweighted, and individuals are less likely to purchase insurance. In the language of optimism and pessimism, individuals tend to be pessimistic about risks that involve large, unlikely losses, but are optimistic about risks that involve large, likely losses.

The above argument is independent of the shape of the utility function. Most psychologists and economists believe that the utility function is S-shaped, being risk-seeking over losses and risk averse over gains (Kahneman and Tversky 1979). If utility is risk-seeking over losses, then an expected utility maximizer would never purchase insurance unless it was heavily subsidized. However, the pessimism of the probability transformation function makes a rank-dependent utility maximizer more risk averse when the probability of the loss is small, and if the pessimism of the probability transformation function is enough to overcome the risk-seeking effect of the utility function, the individual purchases insurance. When the probability of the loss is large, both the utility function and the probability transformation function contribute to risk seeking, and the individual would not purchase insurance against likely risks.

The opposite pattern holds with gambling. People often bet on small chances to win large prizes, but do not bet on large chances to win large prizes. For example, many people would forego a dollar in favor of a 1/10,000 chance of winning $10,000, but not many people would forego $9000 in favor of a 9/10 chance of winning $10,000. The probability transformation function underweights large probabilities of a gain (because they correspond to a small probability of the worse outcome) and overweights small probabilities of a gain. Since people are generally considered to have risk averse utility functions over gains, the contribution of the probability transformation function makes them even more risk averse when the gain is likely but may make them risk seeking when the gain is unlikely.

A second application of the pattern of overweighting unlikely, extreme outcomes arises from the study of sequential search. Consider the problem of an individual searching for the lowest price. Each period that he searches he pays a search cost and then draws a price independently from a random price distribution. The standard economic model, based on expected utility, states that the individual should continue searching as long as the expected marginal benefit from further search exceeds the marginal cost (Rothschild 1974). The expected benefit of search is computed as the additional utility the individual receives from the lower price multiplied by the probability of finding that price, summed over all lower prices. If the individual is a rank-dependent utility maximizer, though, this calculation must be changed, with the transformed probability taking the place of the true probability in determining the marginal benefit of further search. The most extreme gain corresponds to the lowest price, and if finding that lowest price is unlikely, the rank-dependent utility maximizer overweights the gain from the lowest price. This leads to more intensive search with rank-dependent utility. In the language of optimism and pessimism, when the probability of the lowest price is small, the searcher is optimistic about the process, and searches more intensively.

Rank-dependent utility also predicts a greater intensity of research and development activity when the probability of an extremely successful development is small. With the expected utility model, an R&D project is undertaken when the expected benefit from the project exceeds the cost. With rank-dependent utility, the expectation of the benefit is taken using the transformed probability distribution, which exaggerates success when it is unlikely and understresses failure when it is likely.

The results for search and for R&D activity have the same flavor. Rank-dependent utility maximizers spend more resources trying for the big prize than an expected utility maximizer would. This logic should extend to contests in which participants expend effort in the hopes of winning a prize. If the number of contestants is large, the probability of winning should be suitably small. In these cases, rank-dependent utility overweights the small probability of winning, leading contestants to exert more effort in an attempt to win than expected utility theory would predict.

A third application of the rank-dependent utility pattern of overweighting unlikely, extreme events, this time with policy implications, arises from the study of crime and punishment. An individual commits a crime, resulting in a gain. With some probability the offender is caught and punished, resulting in a loss. An expected utility maximizer weighs the gain from crime against the expected punishment when deciding whether or not to commit the crime (Becker 1968). A rank-dependent utility maximizer transforms the probability of punishment in his benefit-cost calculations, and whether he underweights or overweights the punishment probability depends on how high that probability is. If the punishment probability is low, the individual overweights the probability, and criminal behavior is deterred more than the expected utility model predicts. If, on the other hand, the punishment probability is high, the individual underweights the probability, and deterrence is reduced compared to the expected utility case.

Policy implications arise because the law enforcement agency can, at least in part, choose the probability of punishment. In some cases the probability of punishment is very low. For example, Bernasconi (1998) analyzes tax compliance in the United States. The probability of being audited in the U.S. is, on average, less than 2%. This means that the probability of punishment is in the range that is overweighted by rank-dependent utility maximizers. Because of the overweighting, rank-dependent utility predicts less cheating than expected utility does.

Given that rank-dependent utility maximizers overweight small probabilities of bad outcomes, and given that law enforcement is costly, agencies that set low punishment probabilities are getting more "bang for the buck" than agencies that set high punishment probabilities. When the probability is low the decision weight attached to punishment is enhanced, but when the punishment probability is high the decision weight attached to punishment is dampened. Standard models of costs suggest that law enforcement becomes increasingly costly as the probability of punishment rises, implying that punishment probabilities should be kept low enough to allow the probability transformation function to enhance the probability.

All of these applications show that the overweighting of extreme, unlikely events has potentially important implications for both understanding and predicting behavior and also for policy making. But, the actual relevance of these implications is yet to be established. The evidence for the overweighting of extreme, unlikely events is still incomplete. Experimental subjects seem to show the pattern in pairwise choice tasks. The fact that people insure against unlikely losses and gamble on unlikely gains is clearly consistent with the pattern, but may arise from other considerations, and is certainly not direct evidence of the shape of the probability transformation function. Before the implications of the shape of the probability transformation function should be used to form policy, the shape should be investigated using real-world data on real decisions in a variety of contexts. Only then will policy-makers have sufficient confidence in the underlying assumptions of the model to warrant basing policy on them.

5. IMPLICATIONS II – INERTIA AROUND CERTAINTY

The probability transformation function is not the only feature of rank-dependent utility that makes it different from expected utility. The rank-dependent functional form itself also has implications that are separate from the implications of the shape of the transformation function. The feature of rank- dependent utility exploited here is the fact that the outcomes of the distribution are ordered from lowest to highest before the probability transformation is applied, and the property that arises is inertia around certainty.[5]

Consider a simple portfolio allocation problem. There are two assets in which an individual can invest. If he invests everything in the safe asset, his portfolio will be worth x in a year. If he invests everything in the risky asset, in a year his portfolio will be worth y with probability p and z with probability $1 - p$. To make the problem interesting, assume that $0 < y < x < z$, so that neither asset obviously dominates the other. The investor must choose what fraction a of his wealth to invest in the risky asset. If $0 < a < 1$ his portfolio contains a mixture of the two assets. If $a > 1$ he purchases a negative amount of the safe asset, which can be interpreted as borrowing. If $a < 0$ he purchases a negative amount of the risky asset, which can be thought of as selling short.

An expected utility investor chooses a to maximize

$$EU(a) = pu(ay + (1 - a)x) + (1 - p)u(az + (1 - a)x). \tag{22}$$

The optimal value of a can be found by setting the derivative of $EU(a)$ equal to zero:

$$pu'(ay+(1-a)x)[y-x]+(1-p)u'(az+(1-a)x)[z-x] = 0. \tag{23}$$

If u is concave, which it tends to be when outcomes are gains, the second order condition for maximization is satisfied.

We are interested in circumstances in which the investor chooses not to take a position in the risky asset; that is, situations in which $a = 0$. A necessary condition for this can be found by fixing x, y, and p, and setting $a = 0$ in Equation (23):

$$pu'(x)[y + x] + (1 - p)u'(x)[z - x] = 0. \tag{24}$$

Under the standard assumption that $u'(x) > 0$ for all x, this expression can be solved for $z_0 = (x - py)/(1 - p)$. When $z > z_0$ the risky asset becomes more attractive and he invests a positive amount in it, and when

$z < z_0$ the risky asset is less attractive and he invests a negative amount in it. The important point, though, is that there is only one value of z for which the investor takes no position in the risky asset.

Now consider the case of a rank-dependent utility investor. Remember that it now matters which of the payoffs is larger, which in turn depends on whether the investor buys the risky asset or sells it short. When he invests a in the risky asset the outcomes are $ay + (1-a)x$ and $az + (1-a)x$ with $y < x < z$, so when $a > 0$ the lower outcome is $ay + (1-a)x$, and when $a < 0$ the lower outcome is $az + (1-a)x$. For positive values of a, then, the rank-dependent utility is

$$V(a) = h(p)u(ay + (1-a)x) + [1 - h(p)](az + (1-a)x), \tag{25}$$

and for negative values of a, the rank-dependent utility is

$$V(a) = [1-h(1-p)]u(ay+(1-a)x)+h(1-p)(az+(1-a)x). \tag{26}$$

Now follow the same logic as in the expected utility case. The investor chooses $a > 0$ when the derivative of (25) is positive at $a = 0$, and chooses $a < 0$ when the derivative of (26) is negative at $a = 0$. The former condition reduces to $z > [x - h(p)y]/[1 - h(p)]$, and the latter reduces to $z < [x-[1-h(1-p)]y]/h(1-p)$. The rank-dependent utility investor takes no position in the risky asset if

$$\frac{x - [1 - h(1-p)]y}{h(1-p)} \leq z \leq \frac{x - h(p)y}{1 - h(p)}. \tag{27}$$

In the expected utility case, $h(1-p) = 1 - h(p)$, so the left- and right-hand terms of (27) are identical and equal to $(x - py)/(1 - p)$, as above. For general rank-dependent preferences, however, the two are unlikely to be equal. If the left-hand term is less than the right-hand term, (27) defines a range of values of z for which the investor takes no position in the risky asset, that is, he neither buys it nor sells it short.

This behavior can be interpreted as inertia around certainty. With expected utility, certainty is a knife-edge condition; it occurs only for a single, precise value of z. This value of z is easy to interpret. If $z_0 = (x - py)/(1 - p)$ then $py + (1 - p)z_0 = x$, so the risky and safe assets have the same expected payoff. The risky asset is risky, though, so a risk averse investor avoids it. For any other value of z, however, the two expected payoffs are different, and the expected utility investor takes a nonzero position in the risky asset. So, any infinitesimal move away from z_0 moves the investor away from certainty. With rank-dependent expected utility, certainty is no longer a knife-edge condition. Once a value of z is found in the interior of the interval in (27), it takes a discrete (non-infinitesimal) change in z to induce the investor to move away from certainty.

Inertia around certainty occurs when $h(1 - p) > 1 - h(p)$, which deserves more exploration. Rearranging this condition yields $h(p)+h(1-p) > 1$, which can occur when the function h is subadditive.[6] To understand how it works, look back at the investor's decision. When deciding to buy the risky asset, the low outcome occurs when the asset takes the low value y, and this outcome is weighted by $h(p)$ in Equation (25). When deciding to sell the risky asset short, the low outcome occurs when it takes the high value z, and this outcome is weighted by $h(1 - p)$ in Equation (26). Subadditivity implies that together the two of these weights are large, or at least larger than they would be under expected utility. When $h(p)$ is large, buying the risky asset is unattractive, and when $h(1 - p)$ is large, selling it short is also unattractive. So, the investor does neither.

It is also possible that $h(p) + h(1 - p) < 1$, consistent with h being superadditive. In this case the investor prefers *both* buying the risky asset and selling it short to taking no position. Consequently, when $h(p)+h(1-p) < 1$, there is a different sort of inertia around certainty. This time, though, the investor's position in the risky asset is bounded away from zero.

The probability transformation function given in expression (20) and shown in Figure 1 is superadditive, not subadditive. But, casual observance of asset markets finds that investors take positions in only a small minority assets, which is inconsistent with the superadditive function shown in Figure 1. The functional form has also been criticized on other grounds by Neilson and Stowe (2002). While the basic shape of the transformation function seems to fit, with it overweighting low probabilities of the low outcome and underweighting high probabilities of the low outcome, the specific functional form in (20) fails. A more appropriate probability transformation function would have $h(p) + h(1 - p) > 1$ for some values of p. Given that the crossover point is less than $\frac{1}{2}$, though, this inequality will only hold for probabilities outside of the middle range.

Inertia around certainty also pertains in other settings. For example, an expected utility maximizer purchases insurance with full coverage only when the premium is actuarially fair; that is, when the premium is equal to the expected benefit payment received from the insurance company in the case of a loss. If a rank-dependent utility maximizer has a subadditive probability transformation function over the relevant range, he will buy full insurance even if the premium is actuarially slightly unfavorable. In reality, insurance companies charge premia that are higher than their expected payouts, so expected utility would never predict the purchase of full coverage. Many people buy auto insurance with no deductible, though, which is consistent with rank-dependent utility but not with expected utility.

In general, inertia around certainty can be thought of as saying that rank-dependent utility maximizers are more cautious, or more conservative, than their expected utility counterparts. They find the safe alternative attractive under a wider variety of circumstances, and small changes in the environment are less apt to make them change their position away from the safe one. This is related to, but fundamentally different from, risk aversion. Expected utility maximizers can be risk averse, but for them certainty is always a knife-edge condition. For rank-dependent utility maximizers, however, certainty holds a special attraction (or repulsion) that other probability distributions do not possess.

6. CONCLUSION

The rank-dependent utility model, with its probability transformation function, represents an expansion in flexibility over the standard expected utility model, and it gives rise to new notions of attitudes toward risk and to two new behavioral patterns that promise to be of importance. The new notions are optimism and pessimism, and the two new patterns are the overweighting of unlikely, extreme outcomes and inertia around certainty. The fact remains, though, that expected utility is the foundation of rational choice theory, and the decision weights it assigns to outcomes are the probabilities of those outcomes. In rank-dependent utility, because of the probability transformation function, the decision weights differ from the probabilities. Can rank-dependent utility be considered rational, then, or is rank-dependent utility merely a means of describing a certain type of irrationality?

So far in this paper we have addressed three distinct "layers" of "rationality". One is that preferences are complete and transitive. Letting $\succeq$ denote the relation "preferred or indifferent to", the preference ordering is complete if for any p and q in P, either $p \succeq q$, or $q \succeq p$, or both. It is transitive if for any p, q, r in P, $p \succeq q$ and $q \succeq r$ implies that $p \succeq r$. Clearly if a function $V\colon P \to \mathcal{R}$ represents preferences in the sense that $p \succeq q$ if and only if $V(p) \geq V(q)$, these two conditions must hold by the property that the real numbers can be ordered. More specifically, completeness must be satisfied because V will rank any two lotteries according to which is the greater of $V(p)$ and $V(q)$, and V will exhibit transitivity because if $V(p) \geq V(q)$ and $V(q) \geq V(r)$, then $V(p) \geq V(r)$. All of the preferences discussed above, including expected utility, rank-dependent utility, and prospect theory, satisfy completeness and transitivity, and so exhibit the first layer of rationality.

The second layer of rationality requires that preferences at least be capable of exhibiting first-order stochastic dominance preference. First-order stochastic dominance preference is consistent with the notion that the individual likes movements to unambiguously better probability distributions, as in Section 2, and is the correct stochastic notion of "more is better". As demonstrated in Section 2, both rank-dependent utility preferences and expected utility preferences are capable of exhibiting this property, but prospect theory, which transforms each probability individually, is not.

The third layer of rationality is that the preference function use the probabilities of the outcomes as the decision weights on those outcomes. This property is exhibited by expected utility, but not by rank-dependent utility. But, is this layer of rationality as fundamental as the other two, or is it simply a by-product of our familiarity with expected utility as the basis of rational choice theory?

Most economists who are familiar with rank-dependent utility would agree that the first two layers of rationality are fundamental, but that the third is not. A failure of completeness means that the preference ordering is not well-specified, and that upon sufficient reflection an individual should be able to rank two alternatives, even if that ranking is indifference. A failure of transitivity means that the individual can be subjected to a Dutch book process, and can lose money because of it. First-order stochastic dominance preference is an extension of the notion that individuals like money and prefer more of it to less, and any model that must violate first-order stochastic dominance preference is unable to represent the preferences of a selfish person. While it is not necessarily the case that people are always selfish, it is an entirely different matter to preclude selfishness altogether with the preference representation.

Beyond these requirements, everything else is just risk attitude. Any given pairwise choice that is not governed by first-order stochastic dominance preference can be explained by either expected utility or rank-dependent utility as long as the utility function is chosen appropriately. Rank-dependent utility's probability transformation function simply provides a second vehicle for capturing risk attitudes. With expected utility, risk attitudes are governed by the payoffs alone, while in rank-dependent utility risk attitudes are governed by both the payoffs and the probability distribution of the lottery being analyzed. Once one admits a definition of rationality that allows for risk attitudes, why is it necessary to restrict them to being determined by the payoffs and not by the probabilities inherent in the lottery? The fact that rank-dependent utility, through its probability transformation function, can accommodate common combinations of pairwise choices that do not violate first-order stochastic dominance preference but do violate expected utility, simply shows that risk attitudes are too complicated for the expected utility model to handle. The beauty of the rank-dependent utility model is that it captures complex risk attitudes in a compelling, plausible, and useful way.

ACKNOWLEDGMENTS

I am grateful to Ray Dacey and Jack Levy for helpful comments, and to the Private Enterprise Research Center and the Texas A&M Office of Sponsored Projects for financial support.

NOTES

[1] A general proof based on distribution functions is also possible. The distribution F first-order stochastically dominates G if and only if every expected utility maximizer with an increasing utility function prefers F to G, which in turn occurs if and only if $F(x) \leq G(x)$ for all x (see, for example, Hadar and Russell, 1969). Since the probability transformation function h is increasing, $h(F(x)) \leq h(G(x))$ for all x, and every expected utility maximizer with an increasing utility function prefers $h(F)$ to $h(G)$. But this is the same as saying that every rank-dependent utility maximizer with probability transformation function h and an increasing utility function prefers F to G.

[2] Of course, there are other ways in which decision weights can be used to allow for the Allais paradox choices in a rank-dependent utility setting. Focusing on the change in the probability of $0 is the most common.

[3] For example, a P-bet might be an 80% chance of winning $4 (or else winning nothing), and a $-bet might be a 20% chance of winning $16 (or nothing).

[4] The issue of framing is similar, and it is covered in the other papers in this issue. See also Tversky and Kahneman (1981).

[5] Segal and Spivak (1990) call this first-order risk aversion.

[6] Subadditivity is the property that the whole is less than the sum of the parts. In this instance, $h(p) + h(1 - p) > h(p + 1 - p) = 1$.

REFERENCES

Abdellaoui, M.: 2000, 'Parameter-Free Elicitation of Utility and Probability Weighting Functions', *Management Science* **46**, 1497–1512.

Arrow, K. J.: 1974, *Essays in the Theory of Risk-Bearing*, North-Holland, Amsterdam.

Battalio, R. C., J. H. Kagel, and K. Jiranyakul: 1990, 'Testing between Alternative Models of Choice under Uncertainty: Some Initial Results', *Journal of Risk and Uncertainty* **3**, 25–50.

Becker, G. S.: 1968, 'Crime and Punishment: An Economic Approach', *Journal of Political Economy* **76**, 169–217.

Bernasconi, M.: 1998, 'Tax Evasion and Orders of Risk Aversion', *Journal of Public Economics* **67**, 123–134.

Bleichrodt, H. and J. L. Pinto: 2000, 'A Parameter-Free Elicitation of the Probability Weighting Function in Medical Decision Analysis', *Management Science* **46**, 1485–1496.

Camerer, C. F.: 1989, 'An Experimental Test of Several Generalized Utility Theories', *Journal of Risk and Uncertainty* **2**, 61–104.

Camerer, C. F. and T. H. Ho: 1994, 'Violations of the betweenness Axiom and Nonlinearity in Probability', *Journal of Risk and Uncertainty* **8**, 167–196.

Conlisk, J.: 1989, 'Three Variants on the Allais Example', *American Economic Review* **79**, 392–407.

Gonzalez, R. and G. Wu: 1999, 'On the Shape of the Probability Weighting Function', *Cognitive Psychology* **38**, 129–166.

Grether, D. M. and C. R. Plott: 1979, 'Economic Theory of Choice and the Preference Reversal Phenomenon', *American Economic Review* **69**, 623–638.

Hadar, J. and W. R. Russell: 1969, 'Rules for Ordering Uncertain Prospects', *American Economic Review* **59**, 25–34.

Handa, J.: 1977, 'Risk, Probabilities, and a New Theory of Cardinal Utility', *Journal of Political Economy* **85**, 97–122.

Harbaugh, W. T., K. Krause, and L. Vesterlund: 2002, 'Risk Attitudes of Children and Adults: Choices over Small and Large Probability Gains and Losses', *Experimental Economics* **5**, 53–84.

Hershey, J. C., H. C. Kunreuther, and P. J. H. Schoemaker: 1982, 'Sources of Bias in Assessment Procedures for Utility Functions', *Management Science* **38**, 936–954.

Kahneman, D. and A. Tversky: 1979, 'Prospect Theory: An Analysis of Decision under Risk', *Econometrica* **47**, 263–291.

Karmarker, U. S.: 1978, 'Subjectively Weighted Utility: A Descriptive Extension of the Expected Utility Model', *Organizational Behavior and Human Performance* **21**, 61–72.

Lichtenstein, S. and P. Slovic: 1971, 'Reversals of Preference between Bids and Choices in Gambling Decisions', *Journal of Experimental Psychology* **89**, 46–55.

Neilson, W. and J. Stowe: 2002, 'A Further Examination of Cumulative Prospect Theory Parameterizations', *Journal of Risk and Uncertainty* **24**, 31–46.

Pratt, J. W.: 1964, 'Risk Aversion in the Small and in the Large', *Econometrica* **32**, 122–136.

Prelec, D.: 1998, 'The Probability Weighting Function', *Econometrica* **66**, 497–527.

Quiggin, J.: 1982, 'A Theory of Anticipated Utility', *Journal of Economic Behavior and Organization* **3**, 323–343.

Rothschild, M.: 1974, 'Searching for the Lowest Price when the Distribution of Prices is Unknown', *Journal of Political Economy* **82**, 689–711.

Schmeidler, D.: 1989, 'Subjective Probability and Expected Utility without Additivity', *Econometrica* **57**, 571–587.

Segal, U. and A. Spivak: 1990, 'First Order versus Second Order Risk Aversion', *Journal of Economic Theory* **51**, 111–125.

Tversky, A. and D. Kahneman: 1981, 'The Framing of Decisions and the Psychology of Choice', *Science* **211**, 453–458.

Tversky, A. and D. Kahneman: 1992, 'Advances in Prospect Theory: Cumulative Representation of Uncertainty', *Journal of Risk and Uncertainty* **5**, 297–323.

Tversky, A., P. Slovic, and D. Kahneman: 1990, ‘The Causes of Preference Reversal’, *American Economic Review* **80**, 204–217.

Wu, G. and R. Gonzalez: 1996, ‘Curvature of the Probability Weighting Function’, *Management Science* **42**, 1676–1690.

Yaari, M. E.: 1987, ‘The Dual Theory of Choice under Risk’, *Econometrica* **55**, 95–115.

W. S. Neilson
Texas A&M University
Department of Economics
TAMU 4228
College Station, TX 77843-4228
U.S.A.
E-mail: w-neilson@tamu.edu