Certain and Uncertain Utility: The Allais Paradox and Five Decision Theory Phenomena*

James Andreoni†
University of California, San Diego
and NBER

Charles Sprenger‡
University of California, San Diego

January 27, 2010

Abstract

In the study of decision making under risk, preferences are assumed to be continuous. We present a model of discontinuous preferences over certain and uncertain outcomes. Using existing parameter estimates for certain and uncertain utility, five important decision theory phenomena are discussed: the certainty effect, experimentally observed probability weighting, the uncertainty effect, extreme experimental risk aversion and quasi-hyperbolic discounting. All five phenomena can be resolved.

JEL classification: D81, D90

Keywords: Allais Paradox, dynamic inconsistency, risk aversion, uncertainty effect, probability weighting

---

*We are grateful for the insightful comments of many colleagues, including Nageeb Ali, David Eil, Guillaume Fréchette, Uri Gneezy, Glenn Harrison, Muriel Niederle, William Nielson, Joel Sobel, Ulrich Schmidt, Lise Vesterlund participants at the Amsterdam Workshop for Behavioral and Experimental Economics, Stanford Institute for Theoretical Economics, seminar attendees at Harvard University and the University of California, San Diego, and graduate class participants at both Stanford University and the University of Pittsburgh. Andreoni also acknowledges the generous support of the National Science Foundation.

†University of California at San Diego, Department of Economics, 9500 Gilman Drive, La Jolla, CA 92093; andreoni@ucsd.edu.

‡University of California at San Diego, Department of Economics, 9500 Gilman Drive, La Jolla, CA 92093; csprenge@ucsd.edu.

1 Introduction

The Allais common consequence and common ratio paradoxes are known in decision theory as the primary departures from expected utility. Their appeal is that even without experimentation they ring true, and with experimentation they are found to be robust. The two paradoxes were proposed by Allais (1953) using the following hypothetical situations:[1]

1. *Common Consequence:*
   Situation A: Certainty of receiving 100 million.
   Situation B: 10% chance of 500 million; 89% chance of 100 million; 1% chance of nothing.

   Situation A': 11% chance of 100 million; 89% chance of nothing.
   Situation B': 10% chance of 500 million; 90% chance of nothing.

2. *Common Ratio:*
   Situation C: Certainty of receiving 100 milion.
   Situation D: 98% chance of 500 million; 2% chance of nothing.

   Situation C': 1% chance of 100 million; 99% chance of nothing.
   Situation D': 0.98% chance of 500 million; 99.02% chance of nothing.

Situations A and B share a common consequence of winning 100 million with probability 0.89. In situations A' and B' this common consequence is removed. Under expected utility, if A is preferred to B, then A' should be preferred to B', as the manipulation is only one of subtracting a common consequence. Situations C and D have a ratio of probabilities of 0.98. Situations C' and D' have a common ratio. Under expected utility if C is preferred to D, then C' should be preferred to D', as the manipulation is only one of dividing by 100.

Despite the predictions of expected utility, the majority of subjects choose A over B, B' over A' , C over D and D' over C' in similar problems (Kahneman and Tversky, 1979).[2] Allais' initial

[1] French francs were originally used as the currency of the paradoxes. No adjustment made for inflation.

[2] Perhaps the most revealing subject was Leonard Savage who chose A over B and B' over A' and concluded that his preferences were 'subtly' in error (Savage, 1954, p. 103). This may have left Paul Samuelson in an uncomfortable situation as he had stated just before, 'I sometimes feel that Savage and I are the only ones in the world who will give a consistent Bernoulli answer to questionnaires of the type that Professor Allais has been circulating' (Samuelson, 1952, p. 678).

motivation for the paradoxes was an intuition that expected utility's independence axiom was 'incompatible with the preference for security in the neighbourhood of certainty' (Allais, 2008, p. 4). Decision theorists have responded to this critique by relaxing the independence axiom and its implication of linearity in probabilities. The most important associated development is cumulative prospect theory with its S-shaped probability weighting scheme (Tversky and Kahneman, 1992; Tversky and Fox, 1995). Why then, would Allais claim to the present that the paradoxes' true thrust is 'generally misunderstood' (Allais, 2008, p. 5).

One potential source of misunderstanding is that a preference for security in the 'neighborhood of certainty' represents only one half of Allais' intuition. Allais also claimed that 'far from certainty', individuals act as expected utility maximizers, valuing a gamble by the mathematical expectation of its utility outcomes (Allais, 1953).[3] Though the argument is vague as to the definitions of 'neighborhood of certainty' and 'far from certainty', such statements are revealing. In this light, the common ratio and common consequence effects read less like a general violation of linearity in probabilities and more like a local violation that appears as any particular outcome becomes close to perfectly certain. Indeed, if the violation is isolated very close to certainty, it may prove useful to represent it as a violation of continuity. Individuals may exhibit discontinuous preferences over certain and uncertain outcomes. This is similar in spirit to the quasi-hyperbolic representation of discounting where preferences are discontinuous over immediate and future outcomes (Strotz, 1956; Phelps and Pollak, 1968; Laibson, 1997; O'Donoghue and Rabin, 1999). Just as a discontinuity at the present provides a simple, tractable method for representing diminishing impatience through time, a discontinuity at certainty provides a simple, tractable way to represent a disproportionate preference for security near to certainty. To motivate our discussion, we consider the following example:

Decision 1:
Situation P: Certainty of receiving 10 million.
Situation Q: 99% chance of 50 million; 1% chance of nothing.

[3] Allais' wording of 'far from certainty' is 'loin de certitude' (Allais, 1953, p.530, authors' translation).

Decision 2:
Situation P': 99% chance of 10 million; 1% chance of nothing.
Situation Q': 98% chance of 50 million; 2% chance of nothing.

Decision 3:
Situation P'': 98% chance of 10 million; 2% chance of nothing.
Situation Q'': 97% chance of 50 million; 3% chance of nothing.

All three situations share a common ratio of probabilities (with only slight rounding) of 0.99. Under expected utility, if P is preferred to Q, then P' should be preferred to Q', and P'' should be preferred to Q''.[4] Introspection suggests that a substantial proportion of individuals will violate expected utility by preferring P to Q and Q' to P'. However, one would not expect individuals who preferred P' to Q' to prefer Q'' to P''. That is, individuals may violate expected utility between decisions 1 and 2, but not between 2 and 3.[5]

In a survey of 134 University of California, San Diego undergraduate subjects, 52 made decision 1, 40 made decision 2, and 42 made decision 3.[6] While 42.3 percent of subjects preferred P to Q, only 22.5 percent preferred P' to Q', indicating a significant violation of expected utility ($z = 1.999,\ p = 0.046$)[7]. However, the 22.5 percent preferring P' to Q' is not significantly different from the 31.0 percent who preferred P'' to Q'', indicating that expected utility violations are less prevalent away from certainty ($z = -0.864,\ p = 0.388$).[8] Conlisk

[4] Rounding error reinforces this prediction as $0.99 > 0.98/0.99 > 0.97/0.98$. If $U(10) > 0.99 \cdot U(50)$, then $U(10) > 0.99 \cdot U(50) > 0.98/0.99 \cdot U(50) > 0.97/0.98 \cdot U(50)$ by monotonicity. That is, if an individual prefers 10 million with certainty in decision 1, then the 10 million should grow more attractive in decision 2 and even more in decision 3.

[5] Importantly, this occurs in a region where probability weighting is believed to be sharply decreasing (Tversky and Kahneman, 1992; Tversky and Fox, 1995; Prelec, 1998), such that small changes in probabilities are associated with large changes in probability weights. Consider a probability weighting function, $\pi(p)$, monotonically increasing and S-shaped. The elegant probability weighting explanation of the common ratio effect is that the ratio of decision weights away from certainty is larger than the ratio of decision weights closer to certainty. This allows for $U(10) > \pi(0.99) \cdot U(50)$ and $\pi(0.99) \cdot U(10) < \pi(0.98) \cdot U(50)$; $v(10) < \pi(0.98)/\pi(0.99) \cdot U(50)$ as $\pi(0.98)/\pi(0.99) > \pi(0.99)$ when $\pi(\cdot)$ is S-shaped. The probability weighting logic is the same in 1) vs 2) and 2) vs 3) as $\pi(0.97)/\pi(0.98) > \pi(0.98)/\pi(0.99) > \pi(0.99)$ for an S-shaped, monotonic probability weighting function.

[6] The numbers are unbalanced because subjects with ID numbers ending in 0, 1, 2, and 3 were asked to make decision 1. Those with ID numbers ending in 4, 5 and 6 were asked to make decision 2 and those with ID numbers ending in 7, 8, and 9 were asked to make decision 3.

[7] The test statistic z corresponds to the null hypothesis, H_0, of equal proportions preferring P/P' across conditions.

[8] A potential counterpoint is that at probability 0.99, all of the 'action' is effectively taken out of the probability weighting function. This interpretation is at odds with evidence on the shape of the probability weighting

(1989) provides a complementary example, dramatically reducing common consequence effects by moving slightly away from certainty.

The above demonstration helps to pin down the Allais intuition. Away from certainty in decisions 2 and 3, individuals behave roughly consistently. However, when forced to compare situations with certainty and uncertainty together in decision 1, individuals exhibit a disproportionate preference for certain outcomes. Such intuition also carries through to experimental studies. Violations of expected utility are found to be substantially less prevalent when all outcomes are uncertain (Camerer, 1992; Harless and Camerer, 1994; Starmer, 2000). Additionally, a growing body of evidence, at odds with both expected utility and probability weighting, suggests that utility at or near certainty may differ from utility away from certainty (Gneezy, List and Wu, 2006; Simonsohn, 2009; Andreoni and Sprenger, 2009b). In this paper, we demonstrate that allowing for discontinuous preferences at certainty can provide a tractable representation of Allais' intuition and explain a variety of decision theory phenomena including the certainty effect, experimentally observed probability weighting, the uncertainty effect, extreme risk aversion over small stakes and quasi-hyperbolic time preferences.

The paper proceeds as follows: Section 2 discusses continuity of preferences over certainty and uncertainty and presents a simple discontinuous model of preferences. Section 3 summarizes several prior studies suggestive of differences between certain and uncertain utility. In Section 4, these results are applied to the five decision theory phenomena mentioned above. Section 5 provides a brief discussion and conclusion.

2 Continuity over Certainty and Uncertainty

Most textbook treatments of expected utility include an axiom for preference continuity.[9] With fixed outcomes, preferences are represented as an ordering over probability distributions, or lotteries. Let $\mathcal{L}$ represent the space of lotteries and let $\succeq$ be a complete, reflexive and

function, but is similar in spirit to the violation in continuity we discuss below. In Section 4.1, we discuss probability weighting results in more detail.

[9]The definition and treatment here is that of Varian (1992) and is provided primarily for reference.

transitive preference ordering.

Definition (Continuity): Preferences are continuous if the sets $\{p \in [0,1] : p \circ x \oplus (1-p) \circ y \succeq z\}$ and $\{p \in [0,1] : p \circ x \oplus (1-p) \circ y \preceq z\}$ are closed for all $x, y, z \in \mathcal{L}$

Assume there exists a best lottery, b, and a worst lottery, w, such that for any lottery $z \in \mathcal{L},\ b \preceq z \preceq w$. For a given lottery z, there exists a probability mixture of w and b that is in the better than set and another probability mixture of w and b that is in the worse than set. This implies that there exists a third mixture, p_z, such that $p_z \circ b \oplus (1-p_z) \circ w \sim z$.

Importantly, continuity implies that for every certain outcome, a degenerate lottery, there exists a probability mixture of the best and worst elements of $\mathcal{L}$, a non-degenerate lottery, that is indifferent. This further implies that certain and uncertain utility must be functionally identical. The proof is by contrapositive. That is, if certain and uncertain utility are not functionally identical then one can construct a case which violates continuity.

Let there be a fixed set of outcomes. Let $v(x)$ represent utility under certainty and $u(x)$ represent utility under uncertainty and assume that both are increasing in x, but allow $u(x)$ to differ from $v(x)$. Additionally, let there be standard expected utility away from certainty such that for any gamble yielding x_i with probability p_i for $i = 1, 2, ..., S$ with $p_i < 1\ \forall\ i$, total utility is represented by $\sum_i^S p_i u(x_i)$.

For example, take the set of outcomes $\{0, 10, 50\}$ and the set of lotteries (p_1, p_2, p_3) such that $p_1 + p_2 + p_3 = 1$ and $p_i \leq 1,\ i = 1, 2, 3$. Consider the degenerate lottery of 10 with certainty, so $(p_1, p_2, p_3) = (0, 1, 0)$. The best lottery is 50 with certainty, $(0, 0, 1)$, and the worst lottery is 0 with certainty, $(1, 0, 0)$. To satisfy continuity, one must find a probability mixture $(a, 1-a),\ 0 \leq a \leq 1$ such that $a \cdot u(50) + (1-a) \cdot u(0) = 1 \cdot v(10)$. Note that the mixture, because it is a non-degenerate lottery, is necessarily evaluated using the uncertain utility function, $u(\cdot)$, and the certain outcome of 10 is evaluated with $v(\cdot)$. The value $a \cdot u(50) + (1-a) \cdot u(0) < u(50)$ as long as $a < 1$ (a true mixture). If $u(50) < v(10)$, continuity is violated. Figure 1 demonstrates the logic.

Figure 1: Violating Continuity with $u(\cdot) \neq v(\cdot)$

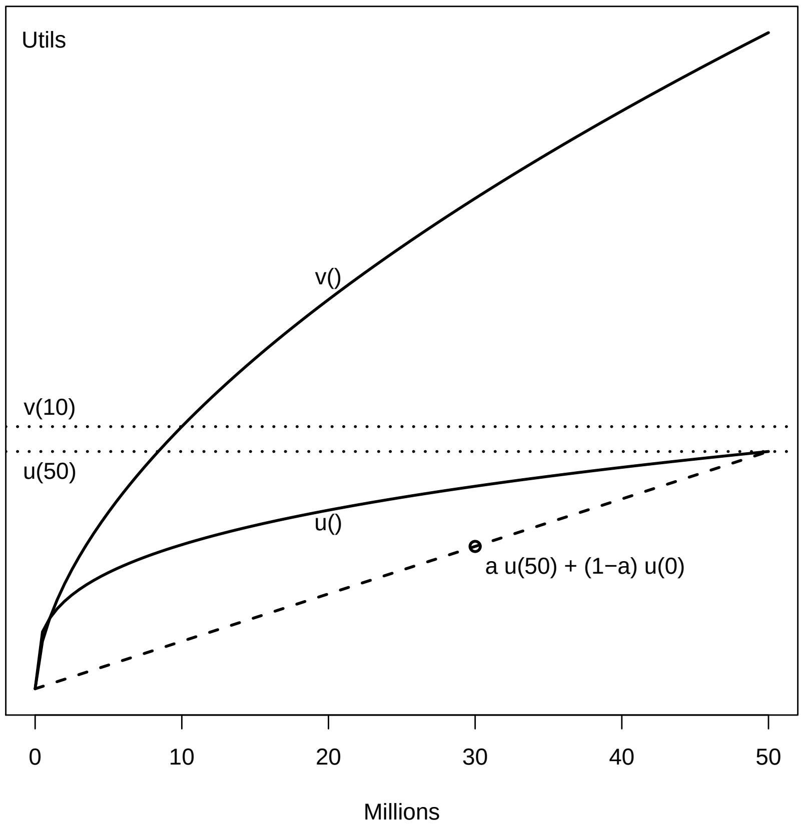

Note: By allowing for a difference between certain utility, $v(\cdot)$, and uncertain utility, $u(\cdot)$, continuity is violated. No probability mixture, $(a, 1-a), a < 1$, exists satisfying $v(10) = a \times u(50) + (1-a) \times u(0)$.

The example discussed above yields exactly the common ratio effect. A disproportionate preference for security at certainty is represented by $v(10) > u(50)$. An individual would choose 10 with certainty over 50 with probability 0.99 if $v(10) > u(50)$. However, an individual would choose 50 with probability 0.98 over 10 with probability 0.99 if $0.98 \cdot u(50) > 0.99 \cdot u(10)$. The common consequence effect can be generated with a very similar argument.

A simple discontinuous utility function can parsimoniously represent the preferences discussed above. Let there be a $1 \times S$ vector of outcomes: $X = (x_1, x_2, ..., x_S)$. Let $\mathcal{L}$ be the set of all lotteries over these outcomes. $\mathcal{L}$ can be partitioned into $\mathcal{L}_D$, the set of degenerate lotteries, and $\mathcal{L}_N$ the set of non-degenerate lotteries. Note that $\mathcal{L}_D$ has exactly S elements, one for each possible degenerate lottery over the S outcomes. Let x_j represent the degenerate lottery outcome associated with a given element of $\mathcal{L}_D$ and let $(p_{N1}, p_{N2}, ..., p_{NS})$ represent a given element of $\mathcal{L}_N$. For a given lottery L, we define the following discontinuous utility function:

$$W(X, L) = \begin{Bmatrix} v(x_j) & if \quad L \in \mathcal{L}_D \\ \sum_{i=1}^{S} p_{Ni} \times u(x_i) & if \quad L \in \mathcal{L}_N \end{Bmatrix}$$

Note that if $u(\cdot) = v(\cdot)$, $W(\cdot)$ reduces to standard expected utility.[10] Continuity is violated if $u(x_i) \neq v(x_i)$ for a given $x_i \in X$. If $v(x) > u(x)$ for $x > 0$, then individuals exhibit a disproportionate preference for certainty.

Critical efforts have been made to provide an axiomatic basis for such a discontinuous representation of preferences (see Neilson, 1992; Schmidt, 1998; Diecidue, Schmidt and Wakker, 2004). These results demonstrate that with additional continuity or independence assumptions, $W(X, L)$ can represent standard expected utility preferences with a disproportionate preference for certainty. However, the represented preferences will violate stochastic dominance (Schmidt, 1998; Diecidue et al., 2004). Though one could work around dominance with an 'editing' argument similar to that of Kahneman and Tversky (1979), this leaves the above model of preferences with the somewhat undesirable property of violating dominance without such additional elements. As such, in the words of Diecidue et al. (2004), 'the interest of the model is descriptive and lies in its psychological plausibility.'

It is important to note that allowing for discontinuous preferences over certainty and uncertainty is not standard in the study of decision making under risk. However, in models of time discounting such preferences frequently arise. Quasi-hyperbolic discounting (Strotz, 1956;

[10]One need not posit standard expected utility as the baseline when $u(\cdot) = v(\cdot)$. Prospect theory probability weighting is also continuous in probabilities. We choose this baseline for two reasons: 1) to have a model with a small deviation from standard expected utility and 2) to capture the stylized fact that away from certainty individuals act mainly in line with expected utility.

Phelps and Pollak, 1968; Laibson, 1997; O'Donoghue and Rabin, 1999) is discontinuous at the present. In such models, individuals discount between the present and one future period with a low discount factor, $\beta\delta$, and between subsequent periods with a higher discount factor, δ alone.

The model of quasi-hyperbolic time preferences is considered an elegant and powerful simplification for several reasons. First, only a single parameter is added to standard exponential discounting; second, many behavioral anomalies can be reconciled. Third, the null hypothesis of dynamically consistent preferences can be tested. Our goal is similar. Consider, for example, standard CRRA utility of $v(x) = x^{\alpha}$. The discontinuity we envision is of the form $u(x) = x^{\alpha-\beta}$, $\alpha > \beta > 0$. A single parameter is added to a standard model, many behavioral anomalies can be reconciled, and the null hypothesis of $\beta = 0$ can be tested. In the next section we show several results in support of modeling such a discontinuity in risk preferences.

3 Experimental Results Suggesting Discontinuity

A growing body of literature is suggestive of discontinuous preferences over certainty and uncertainty. Expected utility is generally found to perform well when all outcomes are uncertain (Conlisk, 1989; Camerer, 1992; Harless and Camerer, 1994; Starmer, 2000).[11] However, violations of expected utility abound when individuals are asked to consider certain and uncertain outcomes together. Generally, the direction of these violations indicate a disproportionate preference for certainty. Of course, these findings alone do not establish differential treatment of certain and uncertain utility. Importantly, there exists a collection of studies which test more directly whether certain and uncertain utility should be treated as identical.

Among the most puzzling and surprising results is the 'uncertainty effect' documented by Gneezy et al. (2006) and reproduced by Simonsohn (2009). In Gneezy et al. (2006), 60 undergraduate subjects at the University of Chicago are randomly assigned to one of three conditions in equal numbers. Subjects were asked to provide their willingness to pay (WTP)

[11] However, there do exist identified violations even when all things are uncertain. For examples and discussion, see the noted citations and Wu and Gonzalez (1996).

for a \$100 gift certificate to a local bookstore, or for a \$50 gift certificate to the bookstore, or for a lottery with 50% chance of winning a \$100 gift certificate and 50% chance of winning a \$50 gift certificate to the bookstore. WTP was elicited using the Becker, DeGroot, Marschak mechanism (Becker, Degroot and Marschak, 1964), 5% of subjects had their choices actualized and were given \$100 to purchase the gift certificate or lottery. The average WTP for the \$50 gift certificate in condition 2 was significantly higher than the WTP for the lottery in condition 3.

The between-subject behavior of valuing a lottery lower than its worst possible outcome violates expected utility and prospect theory probability weighting.[12] The authors suggest that the obtained 'uncertainty effect' should be interpreted as a violation of the 'internality axiom' that 'the value of a risky prospect must lie between the value of that prospect's highest and lowest outcome' (Gneezy et al., 2006, p. 1284). This internality axiom, which has also been called 'betweenness' (Camerer and Ho, 1994) proceeds from continuity. If preferences are continuous, then for a given degenerate lottery, say the \$50 gift certificate, a probability mixture from the strictly better than set and the weakly better than set must be at least as good. As such, violations of the internality axiom could be viewed as violations of continuity. The results of Gneezy et al. (2006) are indicative of certain outcomes being assessed differently than uncertain outcomes. Indeed (Gneezy et al., 2006) suggest subjects may 'code' uncertain lotteries differently than certain outcomes and apply a direct premium for certainty akin to the disproportionate preference $v(x) > u(x)$ for $x > 0$.[13]

In Andreoni and Sprenger (2009b), we present a discounted expected utility violation that is also suggestive of differences between certain and uncertain utility. Using Andreoni and Sprenger (2009a) Convex Time Budgets (CTB) and a within-subject design, 80 subjects are asked to make intertemporal allocation decisions in two principal decision environments. In

[12] Both of these models respect 'betweenness' (Camerer and Ho, 1994), such that lotteries must be valued at some weighted average of the valuations of their outcomes.

[13] The authors express this premium as follows: 'An individual posed with a lottery that involves equal chance at a \$50 and \$100 gift certicate might code this lottery as a \$75 gift certicate plus some risk. She might then assign a value to a \$75 gift certicate (say \$35), and then reduce this amount (to say \$15) to account for the uncertainty' (Gneezy et al., 2006, p. 1291).

Figure 2: Aggregate Behavior Under Certainty and Uncertainty

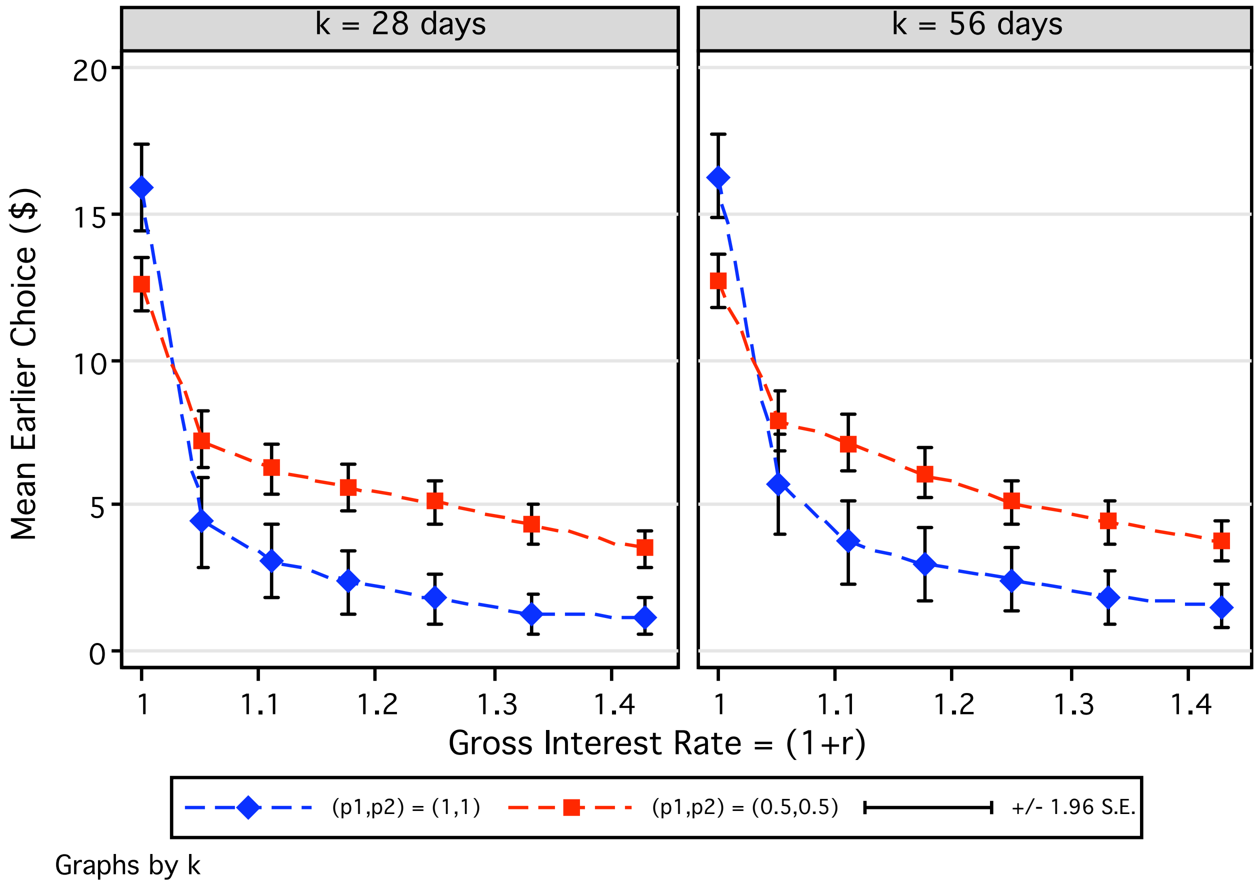

Note: The figure presents aggregate behavior for $N = 80$ subjects under two conditions: $(p_1, p_2) = (1, 1)$, i.e. no risk, in blue; and $(p_1, p_2) = (0.5, 0.5)$, i.e. 50% chance sooner payment would be sent *and* 50% chance later payment would be sent, in red. $t = 7$ *days* in all cases, $k \in \{28, 56\}$ *days*. Error bars represent 95% confidence intervals, taken as $+/- 1.96$ standard errors of the mean. Test of H_0 : Equality across conditions: $F_{14,2212} = 15.66$, $p < .001$.

the first decision environment, sooner and later payments are made 100% of the time.[14] In the second decision environment, sooner payments are made 50% of the time and later payments are made 50% of the time (determined by rolls of two ten-sided die). The prediction from standard discounted expected utility is that allocations should be identical across the two situations. This is due to the common ratio of probabilities across the two condtions. Figure 2 reproduces

[14]See Andreoni and Sprenger (2009b) for efforts made to equate transaction costs.

Figure 2 of Andreoni and Sprenger (2009b), presenting the sooner payment allocation decisions. Allocations in the two decision environments differ dramatically, violating discounted expected utility. Additionally, the pattern of results cannot be explained by either standard probability weighting or temporally dependent probability weighting.[15] In estimates of utility parameters, utility function curvature is found to be markedly more pronounced when all payments are risky as opposed to when all payments are certain, suggesting a disproportionate preference for certainty, $v(x) > u(x)$ for $x > 0$.[16] For experimental payment values of around \$20, certain utility is estimated as $v(x) = x^{\alpha}$, $\hat{\alpha} = 0.988$ ($s.e.$ 0.002), uncertain utility is estimated as $u(x) = x^{\alpha-\beta}$, $\hat{\beta} = 0.105$ (0.017), and the null hypothesis of $\beta = 0$ is rejected.

These studies indicate that modeling utility as different for certain and uncertain outcomes may help to explain decision theory phenomena that remain anomalous in expected utility and probability weighting models. Additionally, many experimental methodologies use certainty equivalence techniques, asking individuals to compare certain and uncertain outcomes. If certain and uncertain utility are different, this may help to explain some of the broad violations of expected utility. In the following section, we discuss five decision theory phenomena existent in the literature that can be explained by allowing certain and uncertain utility to be different. To demonstrate the effects we use the parameter estimates obtained in Andreoni and Sprenger (2009b) in hopes of convincing readers that the difference between certain and uncertain utility need not be large to reconcile anomalous results.

4 Applications: Five Phenomena of Decision Theory

Allowing certain and uncertain utility to be different with a disproportionate preference for certainty can account for five important decision theory phenomena: the certainty effect, experimentally observed probability weighting, the uncertainty effect, experimentally observed extreme risk aversion, and quasi-hyperbolic discounting. Readers will notice that although the

[15]See Andreoni and Sprenger (2009b) for discussion.

[16]Discounting is found to be virtually identical across the two conditions and very similar to Andreoni and Sprenger (2009a).

present section is titled 'Five Phenomena', there are only four subsections. This is because one of the five decision theory phenemona we discuss is trivially generated by a difference between certain and uncertain utility.

The 'certainty effect' is the robust finding, frequently derived from intuitions of the Allais Paradox, that when certain options are available, they are disproportionately preferred. Allowing certain and uncertain utility to differ with the assumption that $v(x) > u(x)$ for $x > 0$ provides the certainty effect trivially. Certain options are assumed to be disproportionately preferred via the functional difference between $u(\cdot)$ and $v(\cdot)$. In what follows, we discuss the other four phenomena in detail.

4.1 Probability Weighting

In addition to the four-fold pattern of risk preferences over gains and losses relative to a reference point, one of prospect theory's major contributions to decision theory is the notion of probability weighting (Tversky and Kahneman, 1992; Tversky and Fox, 1995). Probability weighting generally assumes that there exists a nonlinear function, $\pi(p)$, which maps objective probabilities into subjective decision weights. $\pi(\cdot)$ is normally assumed to be S-shaped. Figure 3 plots the popular one-parameter functional form $\pi(p) = p^{\gamma}/(p^{\gamma}(1-p)^{\gamma})^{1/\gamma}$ with $\gamma = 0.61$ as estimated by (Tversky and Kahneman, 1992).[17] Low probabilities are upweighted and high probabilities are downweighted. Identifying the general shape of the probability weighting function and pinning down its parameter values has received substantial attention both theoretically and in experiments (see e.g., Tversky and Fox, 1995; Wu and Gonzalez, 1996; Prelec, 1998; Gonzalez and Wu, 1999).

Experiments demonstrating an S-shaped probability weighting function generally use certainty equivalence techniques (see Tversky and Kahneman, 1992; Tversky and Fox, 1995). That is, individuals are asked to state a certain amount, C, that makes them indifferent to a lottery which yields X with probability p and 0 otherwise. Analysis proceeds as follows: certain and

[17]Other analyses with similar functional forms yield similar patterns and parameter estimates (for reviews, see Prelec, 1998; Gonzalez and Wu, 1999).

Figure 3: Standard Probability Weighting

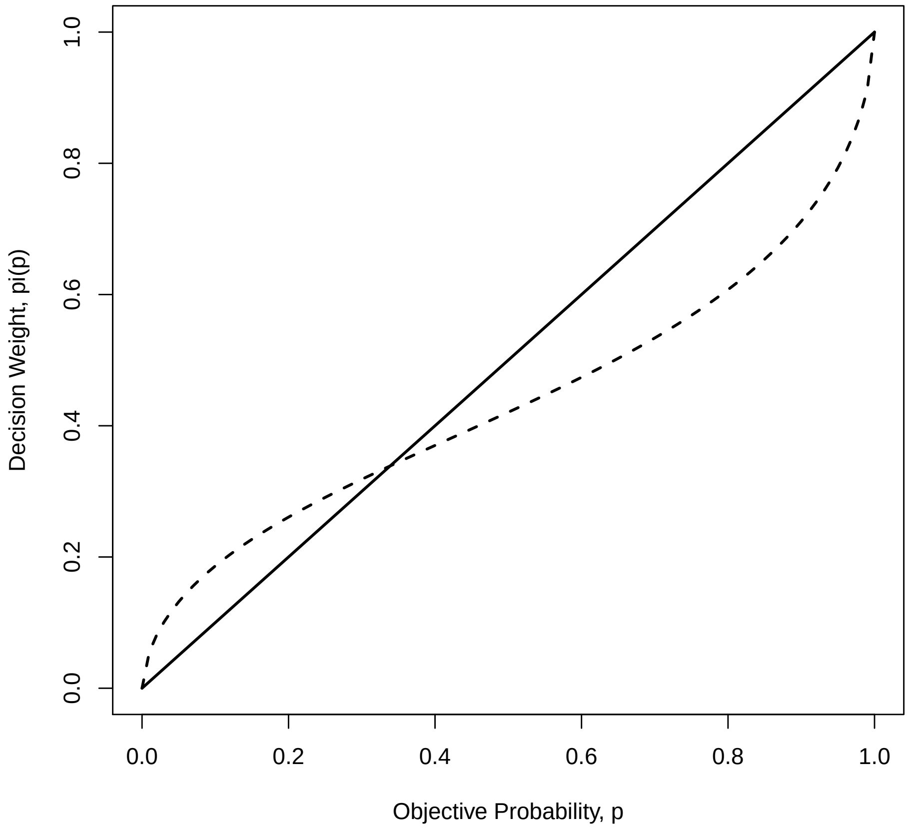

Note: The function $\pi(p) = p^\gamma/(p^\gamma(1-p)^\gamma)^{1/\gamma}$ is plotted with $\gamma = 0.61$ as found by Tversky and Kahneman (1992).

uncertain utility are assumed identical, the utility of zero is normalized to zero, a functional form for utility is posited, and $\pi(p)$ is identified as the value that rationalizes the indifference condition

$$u(C) = \pi(p) \times u(X). \tag{1}$$

In Tversky and Kahneman (1992), $u(X) = X^\alpha$ is posited along with the popular probability weighting function noted above and the parameters of $\pi(\cdot)$ and α are estimated jointly.[18] In Tversky and Fox (1995), the curvature parameter from Tversky and Kahneman (1992) of 0.88

[18]Additional utility parameters such as the degree of loss aversion are also estimated.

is assumed and the parameters of a similar two-parameter $\pi(\cdot)$ function are estimated.

If certain and uncertain utility are functionally different, and $u(0) = 0$, instead of (1) the indifference condition is

$$v(C) = p \times u(X). \tag{2}$$

It is easy to see how this might lead to downweighting of high probability events. Dividing (1) by (2) we obtain $u(C)/v(C) = \pi(p)/p$. If $v(C) > u(C)$, then $\pi(p)/p < 1$. This, of course, assumes that $u(X)$ is parameterized the same way in both cases, which may not be true if one were to estimate the data first assuming (1) and then assuming (2).

Take a gamble with probability $p = 0.90$ of winning \$50. Following the utility parameters obtained in Andreoni and Sprenger (2009b), (2) would yield a certainty equivalent of $C = (0.9 \times 50^{0.99-0.11})^{1/0.99} = \$29.10$. Given a certainty equivalent of \$29.10 and following (1) with the utility parameters of Tversky and Kahneman (1992) and Tversky and Fox (1995), one would infer a decision weight of $\pi(0.9) = 29.10^{0.88}/50^{0.88} = 0.62$, indicating a sharply decreasing probability weighting function away from certainty.[19] Now, take a gamble with probability $p = 0.01$ of winning \$50. Following (2), this would yield a certainty equivalent of $C = (0.01 \times 50^{0.99-0.11})^{1/0.99} = \$0.31$. Given a certainty equivalent of \$0.31 and following (1), one would infer a decision weight of $\pi(p) = 0.31^{0.88}/50^{0.88} = 0.011$, indicating a slight upweighting at low probabilities.[20]

Experiments using certainty equivalence techniques to identify probability weighting would, by our account, suffer from an experimental flaw: the certainty effect is built into the experimental design. Allowing for a difference between certain and uncertain utility, a sharply decreasing probability weighting function and upweighting of low probabilities can be generated. These are the hallmarks of the S-shaped probability weighting function. Interestingly,

[19]It is noteworthy that the curvature of 0.88 from Tversky and Kahneman (1992) is identical to the uncertain curvature parameter of $\hat{\alpha} - \hat{\beta} = 0.88$ from Andreoni and Sprenger (2009b).

[20]The inversion between upweighting and downweighting occurs because under the utility parameterization if $C < 1$, $v(C) < u(C)$ while for $C > 1, v(C) > u(C)$. We recognize that these are likely to be pathological cases, but note that upweighting of low probability events is distinctly less pronounced than downweighting of high probability events (Prelec, 1998).

when the certainty effect is eliminated by design, experimental data appear to reject probability weighting. Andreoni and Harbaugh (2009) eliminate any certain outcomes, ask subjects to trade probability for prize along a linear budget constraint, and find surprising support in favor of expected utility, while significantly rejecting probability weighting.

4.2 The Uncertainty Effect

The uncertainty effect of valuing a lottery lower than its worst possible outcome is a surprising and somewhat unintuitive result. The uncertainty effect violates expected utility, probability weighting and any other utility representation respecting betweenness.

Consider the utility parameters obtained in Andreoni and Sprenger (2009b) and the original uncertainty effect comparing a 50%-50% lottery paying \$50 or \$100 to the certainty of \$50. The utility of the lottery is given as $U_L = 0.5 \times 50^{0.99-0.11} + 0.5 \times 100^{0.99-0.11} = 44.41$. The utility of the certain \$50 is given as $U_C = 50^{0.99} = 48.08$, demonstrating the uncertainty effect of valuing a lottery lower than its worst outcome.

Note should be made of two key issues. First, the uncertainty effect seems not to be present for immediate monetary payments in certainty equivalents experiments (Birnbaum, 1992).[21] Second, it is not observed within individuals (Gneezy et al., 2006). In addition to our hypothesis, future research should examine whether the uncertainty effect is present within individuals and across a variety of rewards, including money.

4.3 Extreme Risk Aversion

Experimentally elicited risk preferences generally yield extreme measures of risk aversion. These results are at odds with standard expected utility theory as even moderate risk aversion over small experimental stakes implies unbelievable risk aversion over large stakes (Rabin, 2000). Importantly, risk preferences are frequently elicited using willingness to pay certainty equivalence techniques similar to the experimental methods used to identify probability weighting.

[21] Though Gneezy et al. (2006) demonstrate that it is present for monetary payments over time.

As such, certainty effects and extreme risk aversion may be conflated.

Allowing for differences between certain and uncertain utility, extreme experimental behavior can be generated. Consider the utility parameters obtained in Andreoni and Sprenger (2009b) and the certainty equivalent of a 50%-50% lottery paying \$50 or \$0. Normalizing $u(0) = 0$, the certainty equivalent is given as $C = (0.5 \times 50^{0.99-0.11})^{1/0.99} = 16.07$.

Under the standard assumption that $u(\cdot) = v(\cdot)$, one would find the curvature parameter, a, which rationalizes $16.07^a = 0.5 \times 50^a$. The solution is $a = 0.61$. Importantly, curvature parameters obtained in low-stakes certainty equivalence studies and in auction experiments are generally between 0.5 and 0.6.[22] This suggests that part of experimentally obtained extreme risk aversion may be associated with differential assessment of certain and uncertain outcomes.[23]

4.4 Quasi-hyperbolic Discounting

In Section 2, we discussed the practice of modeling discontinuous, quasi-hyperbolic time preferences. Recently, arguments have been made that dynamically inconsistent preferences are generated by differential risk on sooner and later payments (for psychological evidence, see Keren and Roelofsma, 1995; Weber and Chapman, 2005). Halevy (2008) argues that differential risk leads to dynamic inconsistency because individuals have a temporally dependent probability weighting function that is convex near certainty similar to standard probability weighting. The probability of receiving payments is argued to decline through time, with present payments being certain. If individuals weight probabilities in a non-linear fashion, then apparent present bias is generated as a certainty effect.

We have demonstrated that if certain and uncertain utility are not identical, certainty effects and sharply declining probability weights can be obtained. As such, one need not call on

[22] In the auction literature, mention is made of 'square root utility' where $\alpha \approx 0.5$. Holt and Laury (2002) discuss several relevant willingness to pay results from the auction literature in line with this value. Interestingly, Kachelmeier and Shehata (1992) present evidence on both willingness to pay and willingness to accept values for lotteries. Though the curvature implied from willingness to pay certainty equivalents is around 0.6, the curvature from willingness to accept treatments actually suggests risk-loving behavior.

[23] Note, our explanation is not sufficient to produce the effect of extreme small stakes risk aversion when comparing two uncertain outcomes as in Holt and Laury (2002). This, in turn, suggests that experimental methodology may also be part of any discussion of extreme risk aversion.

a complex probability weighting function to explain the phenomenon. If individuals exhibit a disproportionate preference for certainty when it is available, then present, certain consumption will be disproportionately favored over future, uncertain consumption. When only uncertain, future consumption is considered, the disproportionate preference for certainty is not active, generating apparent present-biased preference reversals. In Andreoni and Sprenger (2009b), we show that apparent present bias (and future bias) can be generated experimentally via comparisons of certainty and uncertainty.

Consider an individual asked to choose between \$15 with certainty today and \$20 with uncertainty in one month. For simplicity, assume a monthly discount factor of $\delta = 1$ and let $p < 1$ be the assessed probability of being paid in the future. Following the utility parameters of Andreoni and Sprenger (2009b), the relevant comparison is $15^{0.99} = 14.59$ versus $p \cdot 20^{0.99-0.11} = p \cdot 13.96 < 13.96$, and the individual opts for the certain, sooner payment. If asked instead to choose between \$15 with uncertainty in one month and \$20 with equal uncertainty in two months, the comparison is $p \cdot 15^{0.99-0.11} = p \cdot 10.83$ versus $p \cdot 20^{0.99-0.11} = p \cdot 13.96$, the later payment is preferred, and a present-biased preference reversal is observed. This suggests that discontinuous preferences over time and risk may have an identical source: the future is inherently risky.

5 Conclusion

We provide a simple model of discontinuous utility over certainty and uncertainty. We demonstrate that allowing for certain and uncertain utility can help to resolve five decision theory phenomena: the certainty effect, probability weighting, the uncertainty effect, extreme experimental risk aversion, and quasi-hyperbolic discounting. It is compelling that such a broad set of phenomena can be reconciled with discontinuous preferences over certain and uncertain outcomes. Our arguments lay the foundation for re-thinking certain and uncertain utility. It is important to remember that allowing for discontinuous utility need not replace other interpretations of the discussed phenomena, such as visceral present bias or non-linear probability weighting. However, accounting for a disproportionate preference for certainty, following the intuition of Allais, only adds to our understanding of decision-making.

References

Allais, Maurice, "Le Comportement de l'Homme Rationnel devant le Risque: Critique des Postulats et Axiomes de l'Ecole Americaine," *Econometrica*, 1953, *21* (4), 503–546.

_ , "Allais Paradox," in Steven N. Durlauf and Lawrence E. Blume, eds., *The New Palgrave Dictionary of Economics*, 2nd ed., Palgrave Macmillan, 2008.

Andreoni, James and Charles Sprenger, "Estimating Time Preferences with Convex Budgets," *Working Paper*, 2009a.

_ **and** _ , "Risk Preferences Are Not Time Preferences," *Working Paper*, 2009b.

_ **and William Harbaugh**, "Unexpected Utility: Five Experimental Tests of Preferences For Risk," *Working Paper*, 2009.

Becker, Gordon M., Morris H. Degroot, and Jacob Marschak, "Measuring Utility by a Single-Response Sequential Method," *Behavioral Science*, 1964, *9* (3), 226–232.

Birnbaum, Michael H., "Violations of Monotonicity and Contextual Effects in Choice-Based Certainty Equivalents," *Psychological Science*, 1992, *3*, 310–314.

Camerer, Colin F., "Recent Tests of Generalizations of Expected Utility Theory," in Ward Edwards, ed., *Utility: Theories, Measurement, and Applications*, Kluwer: Norwell, MA, 1992, pp. 207–251.

_ **and Teck-Hua Ho**, "Violations of the Betweenness Axiom and Nonlinearity in Probability," *Journal of Risk and Uncertainty*, 1994, *8* (2), 167–196.

Conlisk, John, "Three Variants on the Allais Example," *The American Economic Review*, 1989, *79* (3), 392–407.

Diecidue, Enrico, Ulrich Schmidt, and Peter P. Wakker, "The Utility of Gambling Reconsidered," *Journal of Risk and Uncertainty*, 2004, *29* (3), 241–259.

Gneezy, Uri, John A. List, and George Wu, "The Uncertainty Effect: When a Risky Prospect Is Valued Less Than Its Worst Possible Outcome," *The Quarterly Journal of Economics*, 2006, *121* (4), 1283–1309.

Gonzalez, Richard and George Wu, "On the Shape of the Probability Weighting Function," *Cognitive Psychology*, 1999, *38*, 129–166.

Halevy, Yoram, "Strotz Meets Allais: Diminishing Impatience and the Certainty Effect," *American Economic Review*, 2008, *98* (3), 1145–1162.

Harless, David W. and Colin F. Camerer, "The Predictive Utility of Generalized Expected Utility Theories," *Econometrica*, 1994, *62* (6), 1251–1289.

Holt, Charles A. and Susan K. Laury, "Risk Aversion and Incentive Effects," *The American Economic Review*, 2002, *92* (5), 1644–1655.

Kachelmeier, Steven J. and Mahamed Shehata, "Examining Risk Preferences under Highe Monetary Incentives: Experimental Evidence from the People's Republic of China," *American Economic Review*, 1992, *82* (2), 1120–1141.

Kahneman, Daniel and Amos Tversky, "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 1979, *47* (2), 263–291.

Keren, Gideon and Peter Roelofsma, "Immediacy and Certainty in Intertemporal Choice," *Organizational Behavior and Human Decision Making*, 1995, *63* (3), 287–297.

Laibson, David, "Golden Eggs and Hyperbolic Discounting," *Quarterly Journal of Economics*, 1997, *112* (2), 443–477.

Neilson, William S., "Some Mixed Results on Boundary Effects," *Economics Letters*, 1992, *39*, 275–278.

O'Donoghue, Ted and Matthew Rabin, "Doing it Now or Later," *American Economic Review*, 1999, *89* (1), 103–124.

Phelps, Edmund S. and Robert A. Pollak, "On second-best national saving and game-equilibrium growth," *Review of Economic Studies*, 1968, *35*, 185–199.

Prelec, Drazen, "The Probability Weighting Function," *Econometrica*, 1998, *66* (3), 497–527.

Rabin, Matthew, "Risk aversion and expected utility theory: A calibration theorem," *Econometrica*, 2000, *68* (5), 1281–1292.

Samuelson, Paul A., "Probability, Utility, and the Independence Axiom," *Econometrica*, 1952, *20* (4), 670–678.

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

Schmidt, Ulrich, "A Measurement of the Certainty Effect," *Journal of Mathematical Psychology*, 1998, *42* (1), 32–47.

Simonsohn, Uri, "Direct Risk Aversion: Evidence from Risky Prospects Valued Below Their Worst Outcome," *Psychological Science*, 2009, *20* (6), 686–692.

Starmer, Chris, "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice Under Risk," *Journal of Economic Literature*, 2000, *38* (2).

Strotz, Robert H., "Myopia and Inconsistency in Dynamic Utility Maximization," *Review of Economic Studies*, 1956, *23*, 165–180.

Tversky, Amos and Craig R. Fox, "Weighing Risk and Uncertainty," *Psychological Review*, 1995, *102* (2), 269–283.

_ **and Daniel Kahneman**, "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 1992, *5* (4), 297–323.

Varian, Hal R., *Microeconomic Analysis*, 3rd ed., New York: Norton, 1992.

Weber, Bethany J. and Gretchen B. Chapman, "The Combined Effects of Risk and Time on Choice: Does Uncertainty Eliminate the Immediacy Effect? Does Delay Eliminate the Certainty Effect?," *Organizational Behavior and Human Decision Processes*, 2005, *96* (2), 104–118.

Wu, George and Richard Gonzalez, "Curvature of the Probability Weighting Function," *Management Science*, 1996, *42* (12), 1676–1690.