Submitted to Econometrica

On the Origin of Utility, Weighting, and Discounting Functions: How They Get Their Shapes and How to Change Their Shapes

Neil Stewart, Stian Reimers and Adam J. L. Harris

October 14, 2011

Submitted to Econometrica

ON THE ORIGIN OF UTILITY, WEIGHTING, AND DISCOUNTING FUNCTIONS: HOW THEY GET THEIR SHAPES AND HOW TO CHANGE THEIR SHAPES[1] [2]

NEIL STEWART[a], STIAN REIMERS[b] AND ADAM J. L. HARRIS[c]

We present a theoretical account of the shapes of utility, probability weighting, and temporal discounting functions. In an experimental test of the theory, we systematically change the shape of revealed utility, weighting, and discounting functions by manipulating the distribution of monies, probabilities, and delays in the questions used to elicit them. The data demonstrate that there is no stable mapping between attribute values and their subjective equivalents. This is a profound challenge to expected and discounted utility theories, and also their descendants such as prospect theory and hyperbolic discounting theory. These theories simply assert a particular shape for the stable mapping in order to describe choice data. We explain where the shapes comes from and, in describing the mechanism by which people choose, explain why the shape depends on the distribution of gains, losses, risks, and delays in the environment.

KEYWORDS: revealed utility, cardinal utility, ordinal utility, probability weighting, subjective probability, temporal discounting, delay discounting, stable preferences, decision by sampling, range frequency theory, risky choice, decision under risk, intertemporal choice, decision under delay.

[1]We thank James Adelman for valuable help with the logistic regression analysis and Gordon D. A. Brown, Tatiana Kornienko, Graham Loomes, Will Matthews, Dani Navarro-Martinez, Ganna Pogrebna, and Christoph Ungemach for insightful discussion. This research was supported by ESRC grants RES-062-23-0952 and RES-000-22-3339.

[2]Raw data and the R source code for all analyses are available from the first author.

[a]Department of Psychology, University of Warwick, Coventry, CV4 7AL, UK. neil.stewart@warwick.ac.uk; http://www.stewart.warwick.ac.uk

[b]Department of Psychology, City University, Northampton Square, London, EC1V 0HB, UK. stian.reimers@city.ac.uk

[c]Department of Cognitive, Perceptual and Brain Sciences, University College London, 26 Bedford Way, London, WC1H 0AP, UK. adam.harris@ucl.ac.uk

1. INTRODUCTION

Central to our economic behavior are the attributes money, probability, and time. Our representations of these attributes, and our integration of information across these attributes, is thought to determine our economic behavior. The theories of decision under risk and delay generally assume that we transform money, probability, and delay into subjective equivalents, and then integrate information across these equivalents. These transformations are typically modeled using utility (or value), weighting, and discounting functions. In this paper we present a theory of this process which accounts for the previously observed nature of these functions. A prediction of this theory is that these transformations are not stable, and four experimental tests confirm this prediction. We show that by manipulating the distribution of monies, probabilities, and delays in the question set used to elicit utility, weighting, and discounting functions we can systematically change their shape. That is, we show that if you ask different questions the revealed subjective values of given monies, probabilities, and delays can be adjusted, to some extent, at the experimenter's will. We argue that, although it is possible to derive utility functions, subjective probability functions, and temporal discounting functions from behavioral data (such as a series of choices), these psychoeconomic functions have no psychological reality: There are no look up tables that convert from real-world attributes to their subjective equivalents in people's heads. We also argue that, at the individual level, these functions have no economic use: The functions do not provide a parsimonious description of people's choices, because the functions vary with changes to the context that should be trivial.

Below we review the measurement and use of psychoeconomic functions first in decision under risk (Section 2) and then in decision under delay (Section 3). We illustrate how the psychoeconomic functions are shaped to describe the choices that people make. We then review evidence that suggests that the psychoeconomic functions may be influenced by the distribution of attribute values (Section 4). We present a theoretical account of utility, weighting, and discounting functions that explains both why these functions take the shapes they do and also explains how these shapes are influenced by the distribution of attribute values (Section 5). We present an experimental test of the theory (Section 6). Finally, we close with discussion and implications (Section 7).

2. DECISION UNDER RISK

Expected utility is the normative model for risky decision making (von Neumann and Morgenstern, 1947). In expected utility theory, incremental wealth has diminishing marginal utility—the utility of my second $100, for example, is just a little bit lower than the utility of my first $100. To capture this, wealth is translated into utility by a concave function, initially quite steep but then later more flat showing a high sensitivity to initial increases in wealth but then a lower sensitivity to later increases. To choose between risky options, the average or expected utility of each available course of action is calculated and the action maximizing expected utility is selected.

Economists and psychologists have created a large set of models of decision under risk that follow the expected utility framework. Each model adapts utility theory to incorporate some psychological insight based, often, on observations of people's choice behavior. Prospect theory (Kahneman and Tversky, 1979; Tversky and Kahneman, 1992) is perhaps the most famous, but there are many other significant theories (e.g., Birnbaum, 2008; Birnbaum and Chavez, 1997; Busemeyer and Townsend, 1993; Edwards, 1962; Loomes and Sugden, 1982; Quiggin, 1993; Savage, 1954). In prospect theory terminology, a value function converts money into value (the analogue of utility) and a weighting function converts probability into a decision weight. Figure 1 gives example value and weighting functions.

In fact, prospect theory departs slightly from this description. Original prospect theory has an initial series of editing rules designed to prevent the model from

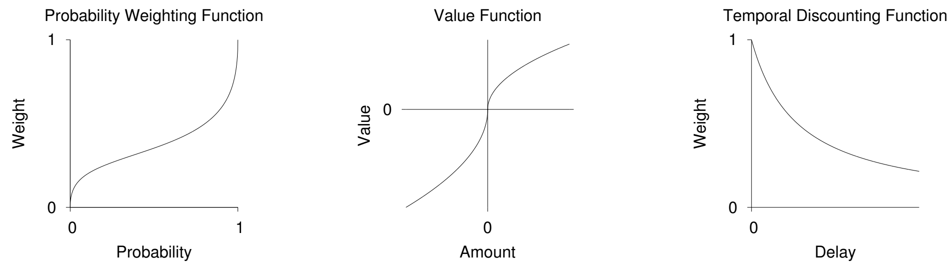

FIGURE 1.— Psychoeconomic functions convert probability (left), money (middle), and delay (right) into their subjective equivalents.

predicting violations of stochastic dominance. Cumulative prospect theory has a slightly different rule for deriving decision weights—cumulative probabilities of achieving at least a given value, rather than raw probabilities, are transformed—also designed to avoid violations of stochastic dominance.

The emphasis in prospect theory was to provide a descriptive account of people's risky decisions, capturing departures from expected utility theory. There are many important empirical departures (see Allais, 1953; Birnbaum, 2008; Camerer, 1995; Loomes, 2010; Luce, 2000; Schoemaker, 1982; Starmer, 2000 for reviews), but here we use just two—the finding that people are risk averse for gains and risk seeking for losses and the common ratio effect—to illustrate how the shapes of the value and weight functions were constructed to account for choice data. Although Kahneman and Tversky did motivate their functional forms by assuming diminishing sensitivity to changes further from reference points (zero for money and zero and one for probability), the value and decision weighting functions are essentially descriptive, motivated to explain particular choice patterns. Consider the shape of the value function for gains (the top right quadrant of the value function in Figure 1. The function is shaped to account for risk aversion in risky choices. For example, people tend to prefer a sure 500 to a 50% chance of 1,000 otherwise nothing. Because the function is concave the value of 500 is a bit more than half the value of 1,000, and so, ignoring probability weighting, the expected value of the sure 500 is higher than the expected value of the 50% chance of 1,000 otherwise nothing. That is, the value function is concave for gains to describe risk aversion. We don't consider losses in this paper, but note here for completeness that the value function is convex for losses because, for moderate probabilities, people tend to be risk seeking for gambles involving losses, rather than risk averse as they are for gambles involving gains (see Tversky and Kahneman, 1992 for a description).

The shape of the weighting function is also constructed to describe the choices that people make. Consider, for example, the common ratio effect (Allais, 1953). In a choice between 3,000 for sure and an 80% chance of 4,000 otherwise nothing, people are risk averse, as above, and prefer the 3,000 for sure. However, in a choice between a 25% chance of 3,000 otherwise nothing and a 20% chance of 4,000 otherwise nothing, people prefer the 20% chance of 4,000 otherwise nothing. Because the second choice is derived from the first by multiplying probabilities by 1/4, participants should prefer the more risky option in both cases or the less risky option in both cases. Thus the empirical finding is not consistent with expected utility theory. Kahneman and Tversky account for these data by assuming that people are most sensitive to changes in small or large probabilities and are least sensitive to changes in intermediate probabilities. Thus the decision weights of 20% and 25% are quite similar but the weights of 80% and 100% are quite different. That is, the weighting function takes its inverse-S shape in order to describe the common ratio effect (and other choice patterns) that people make.

3. DECISION UNDER DELAY

The normative starting point for intertemporal choices is discounted utility theory (Samuelson, 1937). When choosing between a series of delayed outcomes, each outcome is discounted by reducing its value by a constant fraction for every unit of time it is delayed. The value of a delayed outcome thus diminishes exponentially with the length of the delay, and this leads to preferences that remain consistent over time.

As with decision under risk, psychologists and economists have modified this discounted utility framework to incorporate psychological insight based on observation. Perhaps the most significant model is hyperbolic discounting (Ainslie, 1975; Loewenstein and Prelec, 1992, but see also Laibson, 1997; Read, 2001; Scholten and Read, 2010). Here we select one example behavioral finding to illustrate the descriptive model, but there are many (see Loewenstein and Prelec, 1992; Scholten and Read, 2010 for reviews). In the common difference effect people have a tendency to reverse their preferences when a common interval is added or subtracted from the delays of each option. For example, in a choice between 100 in 10 days and 110 in 11 days, people might prefer the later larger reward of 110 in 11 days. However, when the same rewards are brought forwards 9 days, so people are choosing between 100 in 1 day and 110 in 2 days, people might prefer the smaller sooner reward. That is, there is more discounting in moving from 1 to 2 days than from 10 to 11 days (Thaler, 1981). To explain findings like this, Loewenstein and Prelec (1992, see also Mazur, 1987) suggested that discounting was hyperbolic, not exponential (see Figure 1). One property of a hyperbolic function is that the discount rate is initially high and then decreases. Thus large discounting from 1 to 2 days makes the smaller sooner option relatively more attractive but the reduced discounting from 10 to 11 days makes the later larger option relatively more attractive. Note that the motivation for the choice of discounting function is the same as the motivations for the choice of utility and weighting functions: the goal is to describe the choices that people make.

4. MALLEABILITY OF RISKY AND DELAYED CHOICES

Below we review a series of studies that demonstrate that choices and valuations of risky and delayed choices are affected by the distributions of amounts, probabilities, and delays recently experienced. These findings are not consistent with expected utility, discounted utility, or their derivatives. To foreshadow the later theory section, in all of these studies, people behave as if the subjective value of an amount, risk, or delay is given by its rank position in the context created by other recently experienced amounts, risks, and delays. The first set of studies describes how previously encountered amounts, risks, and delays affect current choices. The second set of studies describes how the distribution of potential values affects valuation. The final set describes how choices are affected by the set of alternatives on offer.

In decision under risk, Stewart (2009) demonstrated that a target choice between a 30% chance of 100 points and a 40% chance of 75 points can be reversed by manipulating the distributions of probabilities and amounts encountered just before the choice. When previous choices contained amounts 25, 50, 75, 100, 125, and 150 points and probabilities 30%, 32%, 34%, 36%, 38%, and 40%, people preferred a 40% chance of 75 points. In making this choice, people are behaving as if the difference in amounts is relatively small (the ranks of 75 and 100 points are very similar, ranking 4th and 3rd respectively) but the difference in probabilities is relatively large (the ranks of 30% and 40% are very different, ranking 6th and 1st respectively). If people feel that 40% is much better than 30% but that 100 points is only slightly better than 75 points, a 40% chance of 75 points will be more attractive than a 30% chance of 100 points. In contrast, when previous choices contained amounts 75, 80, 85, 90, 95, and 100 points and probabilities 10%, 20%, 30%, 40%, 50%, and 60%, people preferred a 30% chance of 100 points. In this new context, the ranks of 30% and 40% are very similar but the ranks of 75 and 100 points are very different.

Extending this design, Ungemach, Stewart, and Reimers (2011) have examined how the distribution of attribute values we experience every day affects choices. In one study, Ungemach et al. found that customers leaving a supermarket evaluated the prizes on offer in two simple lotteries against the cost of purchases made a few minutes earlier inside a supermarket. One lottery offered £1.50 with a low probability and the other offered £0.50 with a high probability.

If most purchases were for amounts between £0.50 and £1.50, people behaved as if the difference between £0.50 and £1.50 was larger, and selected the £1.50 lottery. Alternatively, if most purchases were for less than £0.50 or more than £1.50, people behaved as if the difference between £0.50 and £1.50 was smaller, and selected the £0.50 lottery which has the higher probability of winning. Ungemach, Stewart, and Reimers present similar findings for probability when people generate a probability for a weather event before making a risky choice and for delay when people plan for their birthday before making an intertemporal choice. In both this study and the previous study, the previously encountered attribute values were irrelevant to the later choice, but differences in the distributions of the previously encountered attribute values were sufficient to reverse preferences.

The distribution of prices also affects valuations. Birnbaum (1992) and Stewart, Chater, Stott, and Reimers (2003) found that, when valuing a risky gamble, people were influenced by the range and skew of the options available as potential valuations. For example, Birnbaum found that the valuation of a target gamble was higher when people were offered a negatively skewed set of candidate valuations, with more high amounts, than when people were offered a positively skewed set, with more low amounts. Using an incentive compatible auction (Becker, DeGroot, and Marschak, 1964), Ariely, Kőszegi, Mazar, and Shampan'er (2008) found that the distribution of prices from which a sale price was to be randomly drawn affected the reserve price stated by participants in exactly the same way. With negatively skewed sale prices (i.e., more larger prices), stated reserve prices were higher than with positively skewed sale prices (i.e., more smaller prices). In both experiments, a given candidate value or sale price appears larger in the positive-skew condition because there are many smaller prices and few larger ones. In contrast, a given value or sale price appears smaller in the negative-skew condition because there are few smaller prices and many larger ones. Thus, to match the subjective value of target gamble or auctioned good, the candidate value or sale price needs to be larger in the negative skew condition than in the positive skew condition.

In making a single choice, the available options also affect the level of risk demonstrated in that choice. Benartzi and Thaler (2001) examined a natural experiment. Employees were asked to allocate their pension funds between bonds (relatively safe) and stocks (relatively risky). Although all employees were offered at least one stock option and at least one bond option, and thus all employees could exhibit whatever level of risk they preferred, employees made, on average, more risky investments if there were more stock options available. Stewart, Chater, Stott, and Reimers (2003) examine a laboratory analogue, where people are offered either five risky options or five safe options. Participants in the different conditions behaved as though they had different levels of risk aversion, even though participants were randomly assigned to different groups. It is as if the riskiness of an option is a function of its rank position within the context against which it is evaluated.

5. A THEORETICAL ACCOUNT

To summarize the above studies, people behave as if the subjective value of an amount (or probability or delay) is determined, at least in part, by its rank position in the set of values currently in a person's head. So, for example, $10 has a higher subjective value in the set $2, $5, $8, and $15 because it ranks 2nd, but has a lower subjective value in the set $2, $15, $19, and $25 because it ranks 4th.

This suggestion—that subjective value is rank within a sample—is consistent with Parducci's (1965; 1995) range-frequency model of magnitudes. In this model, the subjective value of a magnitude is, in part, given by its rank position. This descriptive model began as an account of scaling of psychophysical quantities, but has more recently been applied in economic contexts such as wage satisfaction (Boyce, 2009; Brown, Gardner, Oswald, and Qian, 2008).

Stewart, Chater, and Brown (2006, see also Stewart, 2009) proposed the decision-by-sampling model of decision making in which the subjective value of an attribute, whether money, probability, or delay, is its rank position in a sample of attributes (the model is motivated by evidence from psychophysical studies of the representation of magnitudes). In decision by sampling, rank emerges from the application of three simple cognitive tools: sampling, binary comparison, and frequency accumulation. Continuing the above example, \$10 is compared to a sample of other amounts in memory: \$2, \$5, \$8, and \$15. In binary comparisons, \$10 looks good compared to \$2, \$5, \$8, but looks bad compared to \$15. By accumulating the number of favorable comparisons across the sample, \$10 is valued at 3/4, because three of the four comparisons are favorable. Stewart, Chater, and Brown (2006) show how these assumptions are sufficient to derive psychoeconomic functions like those in prospect and hyperbolic discounting theories. The left column of Figure 2 shows the distributions of gains, losses, risks, and delays in the real world. The right column gives the resulting utility, weighting, and discounting functions under the assumption that monies, risks, and delays are valued by accumulating the number of favorable comparisons against samples from these real world distributions. For example, because there are more small gains in the world than large gains, the utility function is concave for gains: Against the distribution of gains, a £100 increase in a prize from £100 and £200 improves the rank position substantially more than a £100 increase from £900 to £1,000. Notice the resemblance between these functions and the descriptive functions from prospect and hyperbolic discounting theories in Figure 1. In independent work, Kornienko (2011) shows formally how these tools provide a cognitive basis for cardinal utility. Stewart and Simpson (2008) provide details of a decision-by-sampling process model of Kahneman and Tversky's (1979) data. An on-line decision by sampling calculator, which gives exact choice probabilities for any set of prospects, is available at http://www.stewart.warwick.ac.uk/software/DbS/.

Equation 5.1 formalises the expression for the subjective value of an attribute, although this is scarcely necessary for such a simple theory. The subjective value

$s(x,Y)$ of a target attribute value x in the context of a distribution of n attribute values $Y = \{y_1, y_2, \ldots, y_n\}$ is given by

$$(5.1) \qquad s(x,Y) = \frac{\sum_{y \in Y} c(x,y)}{n}$$

where

$$(5.2) \qquad c(x,y) = \begin{cases} 1 & \text{if } x \text{ compares favorably to } y \\ 0 & \text{if } x \text{ does not compare favorably to } y \end{cases}$$

The subjective value $s(x,Y)$ is the probability of a favorable comparison between x and a randomly selected member of set Y. Equivalently, $s(x,Y)$ is the rank of x in Y, normalized to lie between zero and one.

A strong prediction of these rank-based accounts is that, if one alters the distribution of attribute values encountered, the subjective value of any given attribute should alter too. We call this the *rank hypothesis*. Here we test this prediction for money, probability, and delay. Figure 3 shows hypothetical psychoeconomic functions for the different samples of attribute values used in the present experiments. For each attribute, two distributions are considered. The subjective value of the attribute within the distribution from which it was drawn is plotted against its objective value. Notice how the functions are steepest where the distribution of attribute values is most dense and are shallowest where the distribution of attribute values is least dense. For example, consider the manipulation of the distribution of amounts (middle panel of Figure 3). For the positively-skewed amounts, the value function first increases quickly, as fixed-magnitude increases in amount correspond to large increases in rank within the sample, and then later slowly, as fixed-magnitude increases in amount correspond to small increases in rank within the sample. Thus a positively-skewed set of amounts should give a concave utility function. For the negatively-skewed amounts, the value function first increases slowly, as fixed-magnitude increases in amount correspond to a small increase in rank within the sample, and then later quickly,

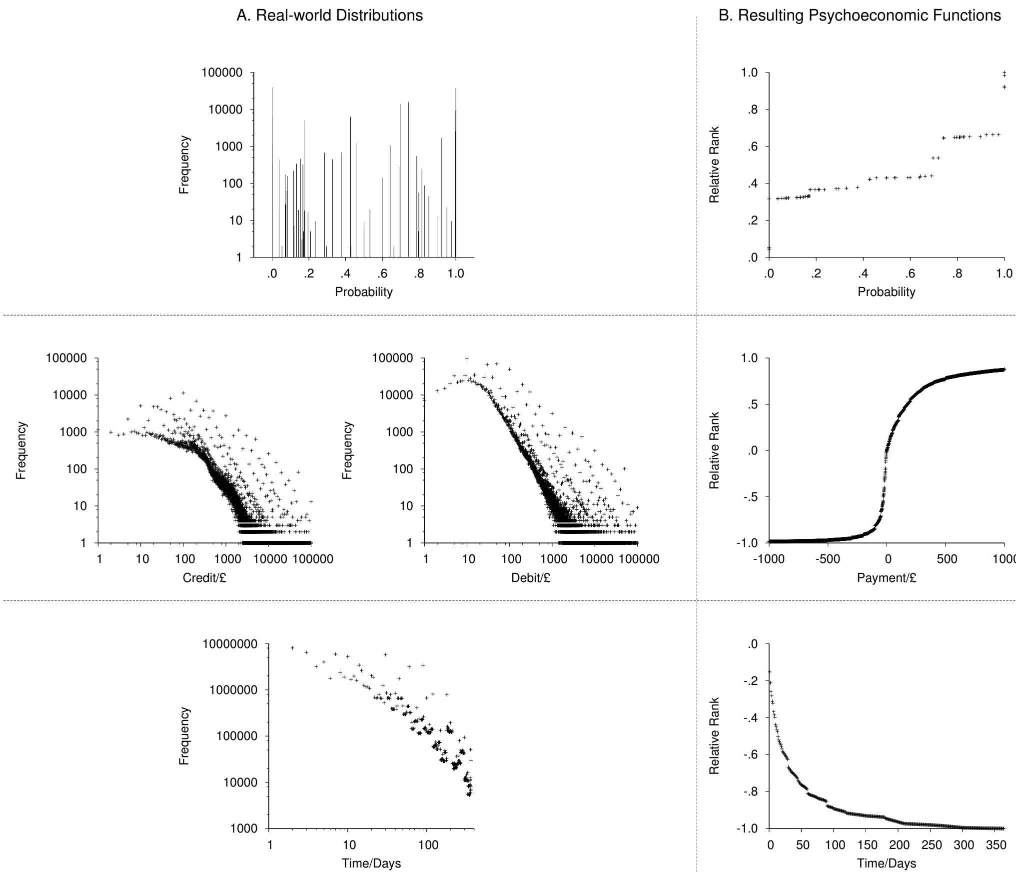

FIGURE 2.— Real world attribute value distributions and the resulting psychoeconomic functions for probability (top), money (middle), and delay (bottom). The panels on the left give the frequency with which different attribute values occur in the real world. Probabilities were phrases (e.g., "possible") used to describe the chances of events sampled from a large corpus of written and spoken language. Gains and losses were sampled from credits and debits to a large sample of checking (in the UK, current) accounts. Delays were sampled from the Internet using Google. The panels on the right give the rank position (normalized between 0 and 1) of an attribute in the sample, and are exactly equivalent to the cumulative frequency distributions. Because larger losses and delays are worse than shorter losses and delays, they have been plotted with negative ranks. After Stewart et al. (2006).

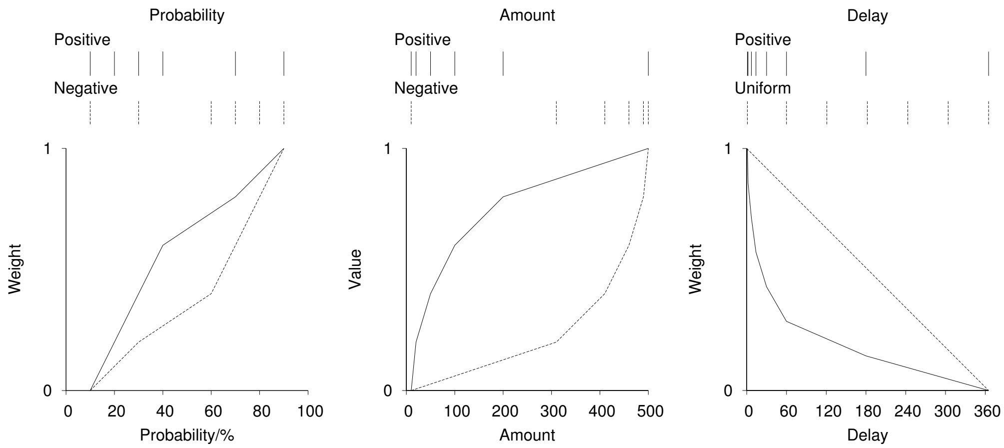

FIGURE 3.— Hypothetical psychoeconomic functions for probability (left), money (middle), and delay (right). The subjective value of an attribute is its rank within the distribution of other attribute values (shown above each plot), normalized to lie between 0 and 1.

a fixed-magnitude increases in amount correspond to larger increases in rank within the sample. Thus a negatively-skewed set of amounts should give a convex utility function. Experiments 1 and 2 explore risky decisions, and investigate how changing the distribution of amounts and probabilities used to build a set of choices changes the psychoeconomic functions revealed from those choices. Experiments 3 and 4 repeat this exercise for intertemporal choices. Experimental payments were incentive compatible and there was no deception.

We do not expect the experimental results to be as extreme as those in Figure 3 because participants are likely to bring previous experience with money, risk, and delay into the laboratory. Consider, for example, the manipulation of the distribution of amounts of money. From Stewart, Chater, and Brown (2006), we know that the distributions of money in the world are positively skewed, with small amounts being highly frequent and larger amounts more rare. Thus in an experimental condition with positively-skewed amounts of money, the distribution the participant has in mind will be a mixture of the positively-skewed experimental distribution and the positively-skewed real-world distribution. Thus the net distribution participants have in mind will be positively skewed and the resulting utility function, under the rank hypothesis, will be concave. But in an experimental condition with a negatively-skewed distribution, the distribution the participant has in mind will be a mixture of the negatively-skewed experimental distribution and the positively-skewed real-world distribution. In this case the net distribution participants experience will be closer to uniform, and thus the resulting utility function, under the rank hypothesis, will be closer to linear. Thus the core prediction is that manipulating the distribution of amounts will give rise to a utility function that is more concave in the positive skew condition compared to the negative skew condition—we predict a relative difference in concavity between conditions, with the absolute concavity being determined (in the decision-by-sampling theory) by the unmeasured contribution of the real-world distribution of amounts. More generally, it is the relative differences in the shapes of the revealed utility, weighting, and discounting functions that test the rank hypothesis, rather than the absolute shapes of these functions.

To preempt later results, manipulating the distribution of amounts, probabilities, and delays alters the pattern of choices that people make, and thus alters the psychoeconomic functions that best describe the data. In fitting these functions, we do not claim that people are using psychoeconomic functions inside their heads. Instead we fit functions to demonstrate that, within the standard frameworks of prospect theory (or subjective expected utility) and of hyperbolic discounting, the shape of a revealed psychoeconomic function is due, at least in part, not to a person's stable risk or intertemporal preferences, but instead to the experimenter's choice of attribute values used in the experiment.

6. EXPERIMENTS

6.1. *Experiment 1*

Participants made a series of 180 choices of the form "p_1 chance of x_1 otherwise nothing" or "p_2 chance of x_2 otherwise nothing". Each choice was between a smaller probability of a larger amount or a larger probability of a smaller amount. Between participants, the distribution of amounts available was manipulated to be either positively skewed (as in the real world, Stewart, Chater, and Brown, 2006) or negatively skewed.

6.1.1. *Method*

Participants

Forty one Warwick psychology first year undergraduates participated for course credit. In addition, for each participant, two gambles were randomly selected to be played for real money. Participants could win up to £5. Data from four participants were deleted for violating stochastic dominance on more than 10% of catch trials (see below), though including these data in the analysis does not alter the pattern of results. Most participants in this experiment, and Experiments 2-4, made no catch-trial errors.

Design

A set of 5 probabilities was crossed with a set of 6 amounts to create 30 gambles of the form "p chance of x". All participants experienced probabilities .2, .4, .6, .8, and 1.0. Participants were randomly assigned to receive either a positively or a negatively skewed set of amounts (see Figure 3, middle panel). The positively skewed set contained amounts £10, £20, £50, £100, £200, and £500. The negatively skewed set was the mirror image of the positively skewed set with the same range, and was constructed by subtracting each amount from £510.

The 30 gambles were crossed with themselves to create a set of choices. Choices between identical gambles and choices where one gamble stochastically dominated the other were dropped to leave 150 choices between a small probability of a large amount and a large probability of a small amount. In addition, 30 of the choices where one gamble stochastically dominated the other were included as catch trials to detect participants who were not making considered choices. All choices are given in Appendix A.

Procedure

Participants were tested individually. Written and spoken instructions explained that participants would be asked to make a series of choices between pairs of gambles. They were told to think of each gamble as an urn draw game in which the urn contained 100 balls, with the percentage of winning balls matching the percentage chance of winning the gamble. It was explained that drawing a winning ball would result in receiving the amount in the gamble and that non-winning balls would result in nothing. Participants were told that they would randomly select two choices at the end of the experiment, with urn draws made to determine their winnings, subject to an experiment exchange rate (which was also applied in the other experiments). The amounts and probabilities on offer were displayed in lists at the top of the screen to remind participants of the attributes they would experience.

Each choice was presented as two buttons, one for each gamble. Each button had text describing the gamble. For example, for one choice in the Positive-Skew Condition, one button was labeled "60% chance of £200" and the other was labeled "100% chance of £10". The assignment of gamble to button was made randomly on each trial. Participants clicked their preferred gamble with the mouse. The next choice appeared automatically. The ordering of choices was set randomly for each participant. A progress bar at the bottom of the screen tracked the progress of the participant through the experiment.

At the end of the experiment, participants randomly selected two choices and played their selected gambles. Participants kept their winnings subject to an experiment exchange rate. Participants could win up to £5.

6.1.2. *Results and Discussion*

For each participant, raw data were 180 choices between pairs of gambles. Two analyses were completed. The first is a nonparametric analysis and reveals utility and weighting functions from the data without assuming a particular functional form for either. The second is a parametric analysis and fits power law utility functions for each participant.

Nonparametric Analysis

Equation 6.1 gives a model for the probability of picking the gamble displayed on the right hand side of the screen as a function of the subjective expected utility $w(p)U(x)$ of each gamble. The weighting function $w(p)$ converts the objective probability p into its subjective equivalent. The utility function $U(x)$ converts money x into its subjective equivalent. Subscripts indicate left and right. The Luce (1959)–Shepard (1957) choice formulation gives a high probability of selecting the right gamble when it has a relatively higher subjective expected utility. The γ parameter controls the degree of determinism in the model: $\gamma = 1$ gives choice probabilities proportional to the subjective expected utilities and $\gamma > 1$ gives more extreme choice probabilities, so gambles with only slightly higher subjective expected utility are very likely to be chosen.

$$P(Right) = \frac{[w(p_R)U(x_R)]^\gamma}{[w(p_L)U(x_L)]^\gamma + [w(p_R)U(x_R)]^\gamma} \tag{6.1}$$

The advantage of using this Luce formulation is that, when expressed as log odds (Equation 6.2), the right hand side of the equation is linear in log subjective values. This means that the subjective values of each probability and amount can be straightforwardly estimated in a logistic regression.

$$\ln\left[\frac{P(right)}{1-P(Right)}\right] = \gamma[\ln w(p_R) + \ln U(x_R) - \ln w(p_L) - \ln U(x_L)] \tag{6.2}$$

Equation 6.3 rewrites Equation 6.2 in matrix notation. Each element in vector β represents the (logarithm of the) subjective probability or value of each attribute used. Each row in matrix $\mathbf{X}$ indicates which attributes are present for a given trial, with -1 marking the presence of an attribute in the left hand gamble, 1 marking the presence of an attribute in the right hand gamble, and 0 the absence of an attribute from that trial. Each element in vector $\mathbf{p}$ represents the probability of responding right, $P(Right)$, for each trial.

$$\ln\left[\frac{\mathbf{p}}{1-\mathbf{p}}\right] = \mathbf{X}.\beta \tag{6.3}$$

Thus standard logistic regression can be used to obtain the maximum likelihood estimate of β. In this way, the subjective probability associated with each probability and the subjective utility of each amount in the experiment are obtained without assuming any functional forms for weighting or utility functions. Without loss of generality, the utility of £500 was set at 1, and the subjective probability of 100% was set at 1.

The revealed utility functions are shown in Figure 4. Effectively, the logistic regression adjusted the heights of each point in the weighting function and utility functions to best fit the choices participants made. Though they were not constrained to do so functions increase monotonically, with the exception of one utility. The utility function for the Positive-Skew Condition is concave, whereas the utility function for the Negative-Skew Condition is convex. Parameter standard errors indicate that the functions differ significantly. And this is the core result: Participants who experienced a different distribution of amounts behaved as if they had a different shaped utility function.

The left panel shows the weighting function, which was assumed to be common to both conditions. The function has a convex shape (cf. the inverse-S shape commonly found, Abdellaoui, 2000; Bleichrodt and Pinto, 2000; Gonzalez and Wu, 1999; Kahneman and Tversky, 1979; Prelec, 1998; Tversky and Kahneman, 1992; Wu and Gonzalez, 1996, 1999). Experimenting with the model fitting shows that there is some degree of trade-off between the shape of the weighting function and the shapes of the utility functions when such simple choices are

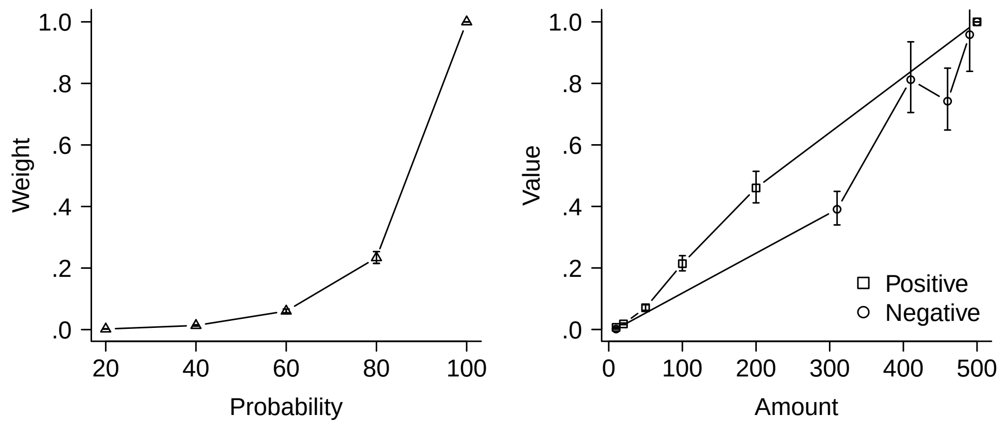

FIGURE 4.— Revealed nonparametric psychoeconomic functions for Experiment 1. Left: The weighting function common to both conditions. Right: The utility functions for the Positive and Negative-Skew Conditions. Error bars indicate the standard errors for the parameter estimates.

used. Because, in our experiment, the size of the amount in a gamble was negatively correlated with the probability of winning that gamble—as is necessary for pairs of gambles where one is not dominated by the other—the relative attractiveness of the risky gamble over the safe gamble can be captured with the shape of the utility function or by the shape of the weighting function. For example, we can make the riskier gamble in a pair more attractive by making the utility function more convex, so the larger amount is relatively more attractive, or the weighting function more concave, so the smaller probability is relatively more attractive. However, because we are interested only in the relative shape of the utility functions across conditions and not the absolute level of concavity, and because the weighting function was common to both conditions, this is not a concern for our analysis. For example, constraining the subjective probability function to be linear increases the concavity of both utility functions. But, the Positive-Skew Condition utility function is always more concave than the Negative-Skew Condition utility function, whether the weighting function is assumed to be concave, convex, or linear.

Do these effects depend on extensive experience with the distributions? We do not think so. Splitting the data for each participant into the first and second halves of the experiment and conducting the analysis separately for each half (for this and later experiments) revealed the same sized effect. This is unsurprising, given attribute distribution effects can take place in as few as 10 trials (Stewart, 2009) or after purchasing a few items in a shop (Ungemach, Stewart, and Reimers, 2011).

Parametric Analysis

The advantage of the nonparametric analysis is that it does not impose a particular functional form on the psychoeconomic functions. A drawback is that data are treated as if all trials come from a single participant. The parametric analysis here has the opposite properties, and imposes a functional form in fitting a power-law function for each participant, and then compares parameters for that power law across conditions. To preempt the results, the same conclusion is reached—that people choose as if they have different utility functions in different conditions.

$$U(x) = x^{\alpha} \tag{6.4}$$

In fitting the data for each participant, utility is assumed to be a power function of amount (Equation 6.4). When $\alpha < 1$, the utility function is concave. When $\alpha = 1$ the utility function is linear. When $\alpha > 1$, the utility function is convex. The choice of a power function is not a theoretical statement from us about the nature of the utility function—it is just a simple function that we can use to fit the data and we would anticipate a very similar result if a different function were used. The weighting function is assumed to be the identity function, so the model reduces to expected utility. Note though that very similar results are obtained if different forms for the weighting function are assumed. The probability of

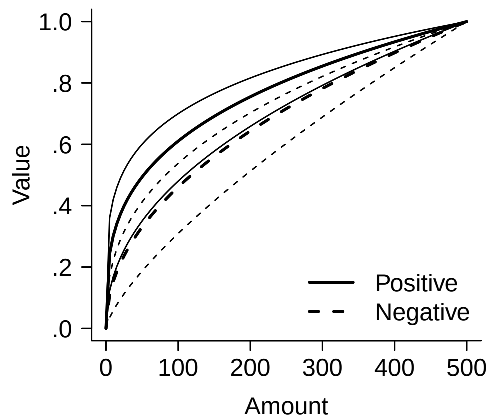

FIGURE 5.— The revealed power-law utility functions for the Positive and Negative-Skew Conditions of Experiment 1. Bold lines are the utility functions of the median participants in each condition. Thin lines indicate the first and third quartile participants.

choosing the right hand gamble is given by Equation 6.1. Here γ and α were estimated by maximum likelihood. Note that very similar results are obtained if γ is fixed at 1.

Figure 5 shows the power utility functions for participants with the median α values in each condition. α values differed significantly between conditions, Welch $t(32.41) = 2.58$, $p = .015$. As in the nonparametric analysis, the utility functions were more concave in the Positive-Skew Condition than the Negative-Skew Condition. The utility functions are both more concave than in the nonparametric analysis (Figure 4) because of the trade-off between the weighting and utility functions in the nonparametric analysis, as described above. When the nonparametric analysis is repeated with a linear weighting function, as was the case in this parametric analysis, the resulting functions are very similar to those found in this parametric analysis. In summary, both the parametric and nonparametric analyses show that the shape of the utility function varied with the distribution of amounts that participants experienced.

6.2. *Experiment 2*

6.2.1. *Method*

In Experiment 2 the distribution of probabilities was manipulated between participants and the distribution of amounts was held constant. In other respects, the method was the same as Experiment 1.

Participants

Thirty five Warwick psychology first year undergraduates participated for course credit. In addition, participants knew they could win up to £5 performance-related pay as in Experiment 1. No participants violated stochastic dominance on more than 10% of catch trials, so all data were retained.

Design

Gambles were made by crossing a set of probabilities with a set of amounts. The amounts were £100, £200, £300, £400, and £500. The set of probabilities was manipulated between participants (see Figure 3, left panel) and was either positively skewed (10%, 20%, 30%, 40%, 70%, 90%) or negatively skewed (10%, 30%, 60%, 70%, 80%, 90%). The negatively skewed set is the mirror image of the positively skewed set. Choices were made by crossing gambles. 120 non-stochastically dominated choices were selected at random and combined with 30 stochastically dominated choices selected at random.

Procedure

Because probabilities were the focus of this experiment, we wanted to be sure that participant understood the probabilities and the method for resolving them. In this study, probabilities were resolved by drawing 1 of 100 chips from a bag. To be successful, a number smaller than or equal to the probability (as a percentage) had to be drawn. For example, for a "70% chance £100" gamble, £100 was received if one of the numbers 1-70 was drawn; numbers 71-100 led to no prize.

This procedure was explained to participants before they commenced the experiment. The experimenter showed participants the chips in an ordered 10 x 10 grid, sweeping their hands over the array to indicate, for several example probabilities, which chips were winning chips. The participant then had the opportunity to ask any questions before the experiment began.

6.2.2. *Results and Discussion*

Nonparametric Analysis

The analysis repeats the procedure from Experiment 1, except subjective probabilities instead of utilities varied between conditions. Figure 6 shows the revealed weighting functions (left). Both functions are convex, but—and this is the crucial result—the Positive-Skew Condition weighting function is less convex than the Negative-Skew Condition weighting function. The probability 70% was common to both conditions, but had a considerably higher subjective weighting in the Positive Skew Condition compared to the Negative Skew Condition. Effectively, people weight outcomes associated with 70% more heavily in the Positive Skew Condition. The right panel shows the utility function common to both conditions to be roughly linear.

Parametric Analysis

The nonparametric analysis suggest that a power law should be an adequate parametric form to describe the weighting function (Equation 6.5). Adapting the analysis for Experiment 1, ϕ and γ were estimated by maximum likelihood, with a linear utility function assumed.

$$w(p) = p^{\phi} \tag{6.5}$$

Figure 7 shows the power weighting functions for participants with the median ϕ values in each condition. ϕ values differed marginally between conditions, Welch $t(26.93) = 1.82$, $p = .080$. As in the nonparametric analysis, the weighting func-

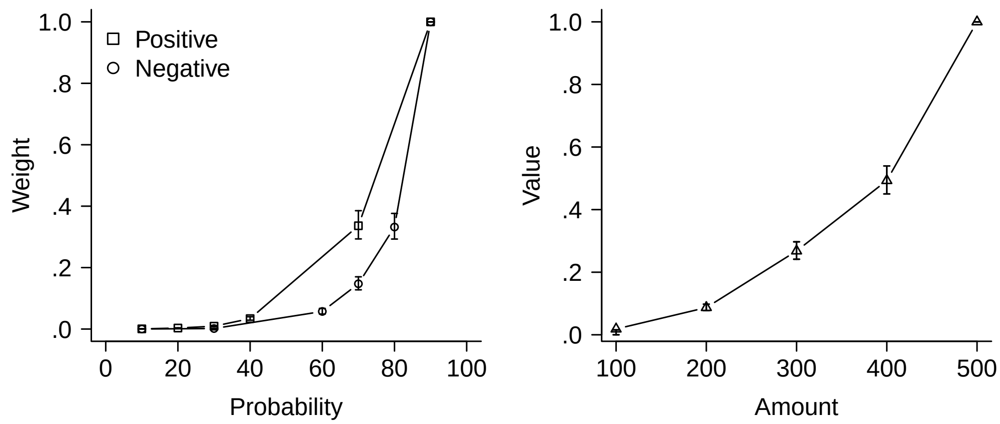

FIGURE 6.— Revealed nonparametric psychoeconomic functions for Experiment 2. Left: The weighting functions for the Positive- and Negative-Skew Conditions. Right: The utility function common to both conditions. Error bars indicate the standard errors for the parameter estimates.

tions were less convex in the Positive-Skew Condition than the Negative-Skew Condition.

6.3. *Experiment 3*

In Experiments 3 and 4 we repeat Experiments 1 and 2, but using delay instead of risk. Participants made choices of the form "x_1 at t_1" or "x_2 at t_2". Each choice was between a smaller sooner amount and a later larger amount. In Experiment 3, the distribution of amounts was manipulated between participants and the distribution of delays was held constant.

6.3.1. *Method*

Participants

Forty Warwick psychology first year undergraduates participated for course credit. Participants knew that one participant would be drawn at random and paid according to one of their choices. Data from four participants were deleted

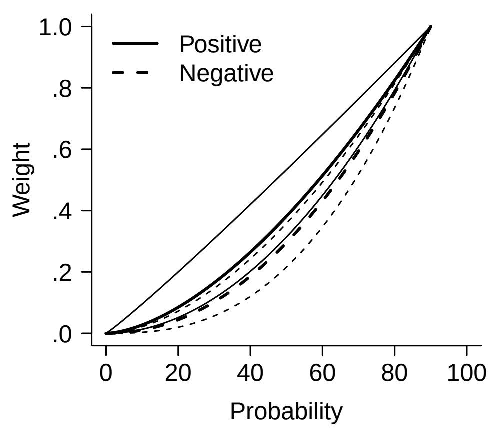

FIGURE 7.— The revealed power-law weighting functions for the Positive- and Negative-Skew Conditions of Experiment 2. Bold lines are the weighting functions of the median participant in each condition. Thin lines indicate the first and third quartile participants. (Note that the almost straight line is the almost linear weighting function for the lower-quartile participant in the Positive-Skew Condition, and is not a benchmark linear weighting function.)

for violating dominance on more than 10% of catch trials, though including these data in the analysis does not alter the pattern of results.

Design

Delayed options were made by crossing a set of delays with a set of amounts. All participants received delays of 1 day, 2 days, 1 week, 2 weeks, 1 month, 2 months, 6 months, and 1 year. The distribution of amounts was either positively or negatively skewed, with values from Experiment 1 (see Figure 3, middle panel). 120 non-dominated choices were selected at random and combined with 30 stochastically dominated choices selected at random.

Procedure

Participants were told that one of them would be selected at random after the experiment, that a choice would be randomly selected, and the winner would be paid according to that choice by bank transfer on the relevant day in the future (after applying an experiment exchange rate). Participants could win up to £20.

6.3.2. *Results and Discussion*

Nonparametric Analysis

The analysis repeats the procedure from Experiment 1, with the subjective values of amounts and times estimated by maximum likelihood. Figure 8 shows the revealed delay discounting and utility functions. The delay discounting function (left) shows a typical hyperbolic-like form. The utility functions (right) differ between conditions. The utility function for the Positive-Skew Condition is relatively linear, whereas the utility function for the Negative-Skew Condition is convex. Crucially, the difference in the utility functions is in the direction expected: the function is more convex in the Negative-Skew Condition than the Positive-Skew Condition.

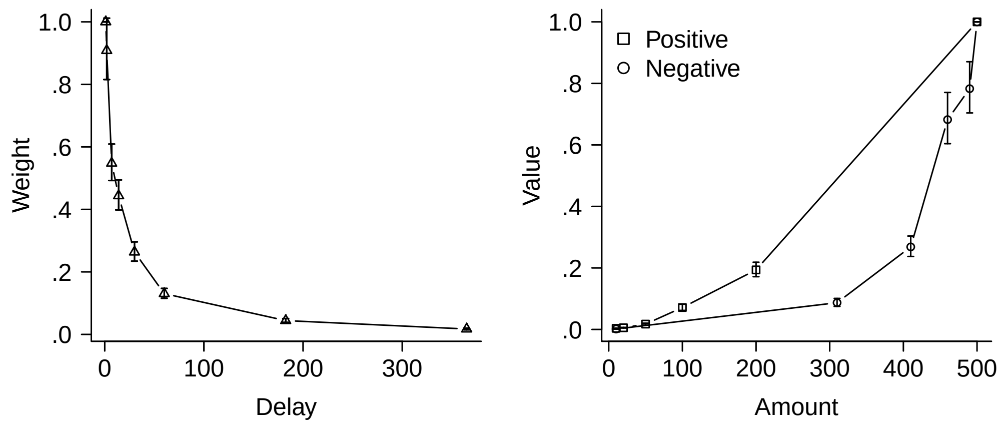

FIGURE 8.— Revealed nonparametric psychoeconomic functions for Experiment 3. Left: The delay discounting function common to both conditions. Right: The utility functions for the Positive-Skew and Uniform Conditions. Error bars indicate the standard errors for the parameter estimates.

Parametric Analysis

As before, utility is assumed to be a power function of amount (Equation 6.4). The delay discounting function was assumed to be hyperbolic, $h(t) = 1/t$. γ and α were estimated by maximum likelihood.

Figure 9 shows the power utility functions for participants with the median α values in each condition. α values differed significantly between conditions, Welch $t(30.01) = 2.53$, $p = .017$. As in the nonparametric analysis, the utility functions were more concave in the Positive-Skew Condition than the Negative-Skew Condition.

In summary, the results of Experiment 3 replicate those of Experiment 1: People behave as if they have different utility functions when the distribution of amounts is different in intertemporal choice as well as in risky choice.

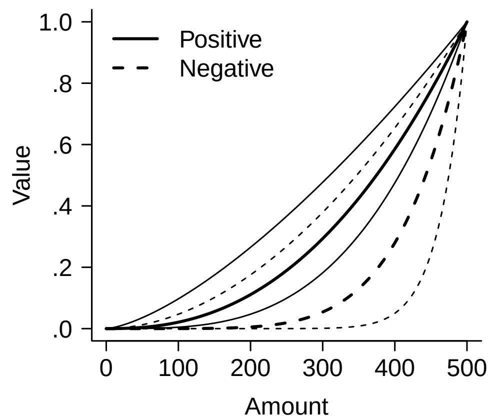

FIGURE 9.— The revealed power-law utility functions for the Positive and Negative-Skew Conditions of Experiment 3. Bold lines are the utility functions of the median participant in each condition. Thin lines indicate the first and third quartile participants.

6.4. *Experiment 4*

6.4.1. *Method*

In Experiment 4 the distribution of delays was manipulated and the distribution of amounts was held constant. In all other respects, the method was the same as Experiment 3.

Participants

Thirty Warwick psychology first year undergraduates participated for course credit. Participants knew that one participant would be drawn at random and paid according to one of their choices. Data from five participants were deleted for violating dominance on more than 10% of catch trials, though including these data in the analysis does not alter the pattern of results.

Design

Gambles were made by crossing a set of probabilities with a set of amounts. The amounts were £100, £200, £300, £400, and £500. The set of delays was manipulated between participants (see Figure 3, right panel) and was either positively skewed (1 day, 2 days, 1 week, 2 weeks, 1 month, 2 months, 6 months, and 1 year) or uniformly distributed (1 day, 2 months, 4 months, 6 months, 8 months, 10 months, and 1 year).

6.4.2. *Results and Discussion*

Nonparametric Analysis

The analysis repeats the procedure from earlier experiments, with weightings of delays allowed to vary between conditions and utilities held constant across conditions. Figure 10 shows the revealed delay discounting functions (left). The delay discounting function is initially much steeper in the Positive-Skew Condition and much closer to linear in the Uniform Condition. For the common 2-month and 6-month delays, people behaved as if they weighted delayed amounts less heavily in the Positive Skew Condition compared to the Uniform Condition. The right panel shows the utility function common to both conditions to be convex.

Parametric Analysis

We use the standard hyperbolic functional form for fitting delay discounting data (Equation 6.6). k and γ were estimated by maximum likelihood, with a linear utility function assumed.

$$h(t) = 1/(1+kt) \tag{6.6}$$

Figure 11 shows the hyperbolic discounting functions for participants with the median k values in each condition. k values differed significantly between conditions, Welch $t(13.93) = 2.21$, $p = .044$. As in the nonparametric analysis, the

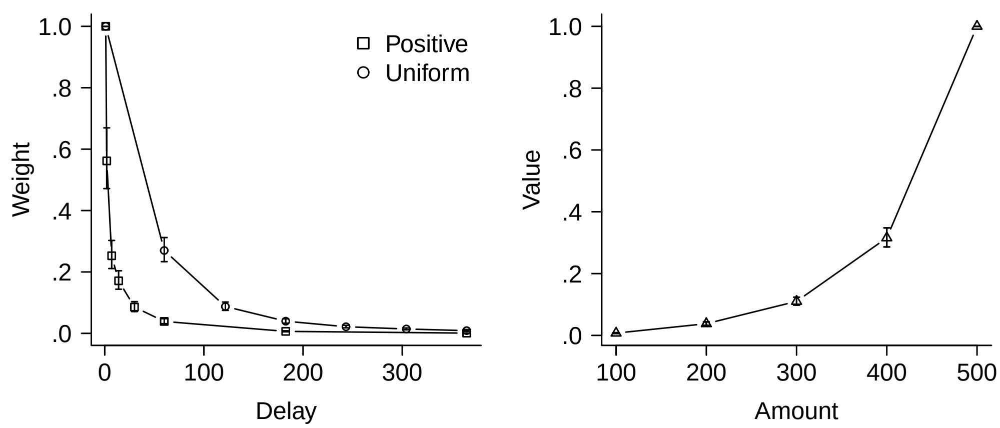

FIGURE 10.— Revealed nonparametric psychoeconomic functions for Experiment 4. Left: The delay discounting functions for the Positive-Skew and Uniform Conditions. Right: The utility function common to both conditions. Error bars indicate the standard errors for the parameter estimates.

delay discounting function was initially steeper in the Positive-Skew Condition and more linear in the Uniform Condition.

7. GENERAL DISCUSSION

Experimentally manipulating the distribution of monies, probabilities, and delays people experience alters the choices people make, which in turn alters the psychoeconomic functions constructed to describe those choices. In Experiments 1 and 3, for risky and intertemporal choices, the revealed utility function translating money into its subjective equivalent was more concave when the distribution of money was positively skewed rather than negatively skewed. In Experiment 2, for risky choices, the revealed weighting function was more concave when the distribution of probabilities was positively skewed rather than negatively skewed. In Experiment 4, for intertemporal choices, the revealed discounting function was more convex when the distribution of delays was positively skewed rather than uniformly distributed. We believe that these are the first experiments

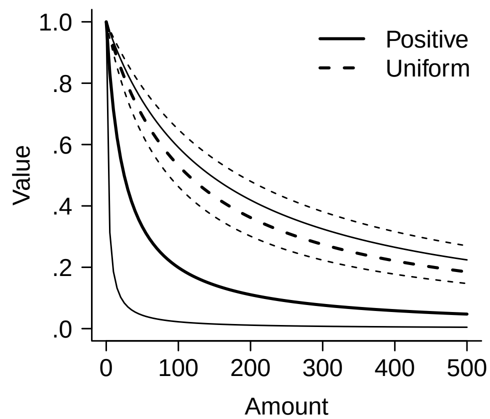

FIGURE 11.— The revealed hyperbolic delay discounting functions for the Positive-Skew and Uniform Conditions of Experiment 4. Bold lines are the delay discounting functions of the median participants in each condition. Thin lines indicate the first and third quartile participants.

to manipulate the shapes of psychoeconomic functions, and are the first experiments to show that the distributions of probabilities and times affect weightings. The effect of the distribution of an attribute value on the subjective function for value was the same for money, probability, and time: The subjective function was steepest when the distribution of attribute values was most dense.

Our data conflict with models which assume that psychoeconomic functions describe a psychological process in which an attribute value is converted into its subjective equivalent. Clearly the psychological processes involved in decision are sensitive to the distribution of attribute values in the immediate context in a way that is not described by the classic psychoeconomic-function approach. A psychological model which assumes an explicit look up, or one in which a look up is used to approximate some other psychological process, does not offer an explanation of the effects we find here. Rather than ignoring effects like these—perhaps treating these effects as noise which could average out by counterbalancing—we believe that they reveal the true way in which decisions are constructed, on-the-fly, using simple cognitive tools.

Sometimes economists are not interested in psychological processes. They present their models as "as if" models, where although people's behavior is consistent with the mathematical model, people are not necessarily assumed to be implementing the calculations in the model. Our data also constrict these "as if" models. Our finding that the distribution of attribute values changes the shape of utility, weighting, and discounting functions limits the descriptive power of the models, because different psychoeconomic functions would be needed for each different choice set. Thus functions revealed in a particular context will not apply in new contexts, reducing the generalizability of the theories. Instead of having a single psychoeconomic function for money, a single psychoeconomic function for probability, and a single psychoeconomic function for delay, a large set of functions would be needed, one for each context. That is, a single function is not a sufficient description of people's choices. Instead a whole set of functions is required, together with a theory to map functions to contexts with different attribute value distributions.

Random utility models (Becker, DeGroot, and Marschak, 1963; see Loomes, 2005, for recent discussion) may seem like an obvious way of accommodating the present results. In a random utility model, the decision maker is assumed to have a set of different utility functions, from which one is randomly selected each time a decision is made (typically, one or more of the parameters of the utility function are assumed to be random variables). These models are intended to account for trial-to-trial variability in the decisions people make (see Blavatskyy and Pogrebna, 2010, for tests of these models). An account of the present results might go as follows: The utility functions drawn could differ across contexts with different attribute-value distributions. What would be required is a theory providing a parsimonious link between the attribute-value distributions in a given context and the set of utility functions available in that context. Such a theory could be descriptive: that is, it could accommodate the effects of attribute-value distributions without explaining how they arise. For example, range-frequency theory might play some role here. But it could go further and explain why the sets of utility functions available differ across contexts with different attribute-value distributions. We offer the decision-by-sampling theory as one candidate here: it suggests that, in the absence of stable internal mappings between attribute values and their subjective equivalents, people are forced to derive subjective values on the fly from a series of comparisons with attribute values they have in mind. In the context of random utility models, the set of available utility functions would be constructed for a particular context, with the variability coming from stochastic components of the memory and comparison processes.

Our experimental results also have implications for the process by which theories of decision under risk and delay are developed. Over the past half century, a considerable literature has been amassed. There are many competing models of decision under risk with, typically, each model accounting for some experimental demonstration of departure from expected utility theory (or prospect theory) or discounted utility theory (or hyperbolic discounting theory). Violations of various axioms of expected or discounted utility theory leads to the development of new models which are based upon alternative axiomatic formulations. Models are also selected on their ability to fit choice data from benchmark data sets. Recently, Brandstätter, Gigerenzer, and Hertwig (2006), Birnbaum (2008), Loomes (2010), and Scholten and Read (2010) have published major theory papers using exactly these methodologies. (Note though that Brandstätter et al.'s model is a set of heuristics, and Loomes's and Scholten and Read's models concerns tradeoffs between attribute values, rather than independent valuation of each option. That is, these models don't follow the standard expected utility or discounted utility frameworks.) Our finding that revealed psychoeconomic functions are not stable presents two problems for this approach. First, axiomatic violations may hold only for certain contexts with particular distributions of attribute values. Of course, any violation rejects a particular model, but an understanding of which distributions lead to which violations will be important in understanding why some models do well in some analyses and badly in others. Second, competitions between models may generalize poorly. Typically competitions involve many choices, and sometimes hundreds of choices, and thus a quite dense and broad distribution of attribute values. Models which perform well in this context may generalize badly to real-world decisions where single choices are made in relatively isolated contexts (apart from in the special case where the distribution of attribute values in the experiment matches the distribution in a participant's memory).

The effects of manipulating the distribution of attribute values are not as complete as would be predicted from the differences in the distributions used in our experiments alone. For example, compare the predicted effect of changing the skew of the distribution of amounts from positive to negative in Figure 3 with the data from Experiments 1 and 3 (Figures 4, 5, 8, and 9). The experimental effects are a little smaller than the predicted effect. One possibility, discussed earlier, is that participants brought to the experiment a background distribution of attribute values from everyday life. This background distribution is probably positively skewed (Stewart, Chater, and Brown, 2006), and would act to dilute the effects of the experimental distributions, making both utility functions more concave. Another possibility is that people do possess some crude underlying transformation between attribute values and their subjective equivalents, but this underlying transformation is extremely malleable. Irrespective of which of these accounts is correct, we advocate abandoning the lookup approach to decision making and we will continue to bring factors affecting the valuation of economic attributes under experimental control.

7.1. *Not Just Another Context / Choice Set Effect*

The literature on choices between gambles documents many context or choice set effects and other violations of expected utility theory. Typically one of two fixes is offered. The first type of fix is a change of psychological primitive. For example, in the endowment effect merely possessing a good makes it more valuable (Kahneman, Knetsch, and Thaler, 1991). The explanation proposed, loss aversion, suggests that people are sensitive to changes (gains and losses) and not final wealth states. As another example, sensitivity to branch splitting suggests people reason with risk-reward branches and not cumulative probability distributions (Birnbaum, 2008). The second type of fix is to assume more complexity in the transformation of objective into subjective values. For example, the four-fold pattern of risk seeking for large losses, risk aversion for small losses, risk aversion for large losses, and risk aversion for large gains and risk seeking for small gains (Kahneman and Tversky, 1979) motivated the inverse-S-shape of the weighting function.

The effects we present here are different. Rather than supporting a change in the shape of a utility, weighting, or discounting function, or a change in the primates which people process, our data suggest that the whole enterprise of using stable functions to translate between objective and subjective values should be abandoned. Theories based around the expected or discounted utility framework are refuted: Descriptively, it is not adequate merely to propose separate functions for each possible context without providing a theory mapping the distribution to the shape of the psychoeconomic function. Psychologically, the translation between objective and subjective values does not represent a process taking place in people's heads.

Is the effect we report large enough that it should be of concern? Our data show that the effect of context here is about as large as the individual differences between people. Figures 5, 7, 9, and 11 show that the difference between the median participants in each condition is about as large as the interquartile range within each condition. If individual differences in, for example, risk aversion are of concern, then we argue that the similar sized effect of context should also be of concern.

7.2. *Previous Investigations of the Shapes of Utility, Weighting, and Discounting Functions*

Here, we illustrate how classic studies in which utility, weighting, and discounting are actually set up with attribute value distributions that produce the classic shapes. The essential observation here is that, quite sensibly, functions are typically measured in more detail (i.e., at more closely spaced intervals) where they are expected to change most quickly. Take, for example, the seminal study by Gonzalez and Wu (1999) on the shapes of utility and weighting functions. Just as we did here, Gonzalez and Wu inferred utility and weighting functions from a series of choices. Gonzalez and Wu used choices between two-branch gambles and sure amounts of money to determine a certainty equivalent for each gamble. They then used a non-parametric method to estimate the utility of each sum of money and the weighting of each probability from these certainty equivalents, just as we did here. Figure 12 plots (open circles) the empirical utility and weighting functions that Gonzalez and Wu recovered. The lines on the plots show the functions predicted by decision by sampling assuming that the set of gambles offered to participants provides the distribution of attribute values against which the amounts and probabilities were compared. For both the utility and weighting functions, the functions predicted under this simple rank hypothesis show the same qualitative pattern seen in the empirical functions recovered by Gonzalez and Wu: The utility function is concave and the weighting function is inverse-S-shaped. Gonzalez and Wu presumably used more small amounts than large amounts because previous research indicated that the utility function varies more quickly over smaller amounts than large amounts so, for accuracy, more measurements should be taken for smaller amounts. Thus there is a self fulfilling result here: By taking more measurements of small amounts, necessarily small amounts occur more often in the question set, which in turn leads to a steeper utility function for smaller amounts. Similarly, Gonzalez and Wu used more small and large probabilities than intermediate probabilities—

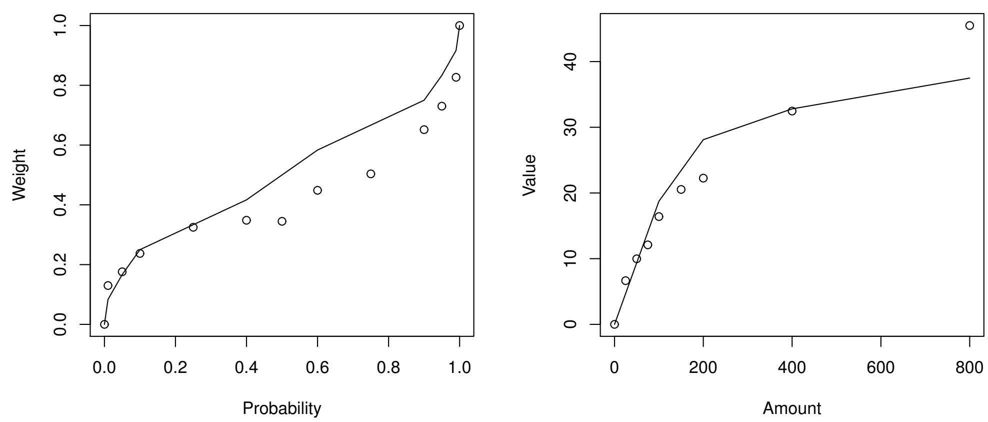

FIGURE 12.— Empirically derived utility (left) and weighting (right) functions from Gonzalez and Wu (1999). Circles plot the functions Gonzalez and Wu recovered. The solid line plots the predictions of decision-by-sampling theory.

because previous research suggests the weighting function is steeper for small and large probabilities and so should be measured more carefully—and thus find a steeper weighting function for small and large probabilities than intermediate probabilities.

The same logic can be applied to classic studies in intertemporal choice. Rachlin, Raineri, and Cross (1991) used hypothetical choices between an immediately available sum of money and a delayed sum. Participants chose between a delayed $1,000 and an immediate sure sum that was titrated up and then down. Then the procedure was repeated for each delay. Figure 13 plots (open circles) the median immediate amount at which people swapped between choosing the delayed $1,000 and the immediate amount. The line on the plot show the function predicted by decision by sampling assuming that the set of delays presented provides the comparison set. As for the Gonzalez and Wu data, the function predicted under this simple rank hypothesis show the same qualitative pattern seen in the empirical function. The discounting function is roughly hyperbolic because of

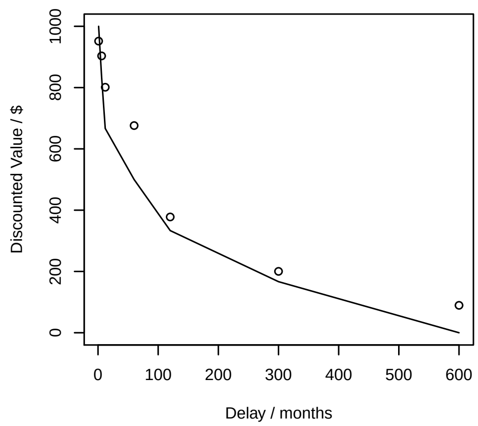

FIGURE 13.— An empirically derived discounting function from Rachlin, Raineri, and Cross (1991), showing the discounted value of a delayed $1,000. Circles plot the function Rachlin, Raineri, and Cross recovered. The solid line plots the predictions of decision-by-sampling theory.

the geometrically spaced distribution of delays chosen by Rachlin, Raineri, and Cross. Virtually all discounting studies that show hyperbolic (or any other kind of) discounting use this kind of distribution of delays.

In summary, we have discussed four findings with one common explanation. The findings are: (a) People behave in the real world as if they have concave utility functions, inverse-S-shaped weighting functions, and hyperbolic discounting functions. (b) People behave in the same way in experiments (as discussed here and in Sections 2 and 3). (c) Changes in the real-world distribution of attribute values cause changes in real world decisions (Section 4, Ungemach, Stewart, and Reimers, 2011). (d) Changes in the distribution of attribute values used in experiments cause changes in laboratory decisions (Experiments 1-4). Expected utility theory and discounted utility theory and their derivatives do not explain any of these results. Here we have argued that all of these findings are the result of the interaction between the distribution of attribute values people have in mind and the simple cognitive tool of binary ordinal comparison described in the decision by sampling theory.

7.3. *Utility*

We finish with a consideration of the implications of our experimental results for the concept of utility.

7.3.1. *Cardinal utility*

Bentham (1970) proposed the idea that the real numbers could be assigned to the pain or pleasure brought by goods. This idea, labeled *cardinal utility*, still remains in risky and intertemporal choice. For example, the idea is central to the development of stochastic models of decision under risk which predict the strength of preference as a function of the difference in utility (e.g., stochastic expected utility theory, Blavatskyy, 2007, and stochastic cumulative prospect theory, Busemeyer, 1985; Erev, Roth, Slonim, and Barron, 2002). Our data present a significant challenge to the cardinal utility approach, because the data are not consistent with a stable mapping between money and the real numbers—the mapping fluctuates between contexts with different distributions of money.

The notion that the subjective value of an attribute might depend on the distribution of attribute values has been incorporated into at least two significant economic models of decision under risk. In Kőszegi and Rabin's (2006; 2007) model, the utility of a particular amount is a combination of classical, consumption utility and gain-loss utility. Gain-loss utility measures the discrepancy between a target amount and the outcomes people expect (see also disappointment theory, Bell, 1985; Gul, 1991; Loomes and Sugden, 1986, where outcomes are evaluated relative to the mean outcome expected rather than to the entire distribution). Specifically, the difference between the target amount and a particular expectation is transformed by a value function that is concave for positive values (where the target amount exceeds the expectation) and convex for negative values (cf. Kahneman and Tversky, 1979), and this transformation is integrated over the distribution of expectations. This model accounts for a series of economic phenomena beyond the expected utility model (e.g., people are willing to pay more for a good that they expected to buy and people are less likely to continue to work when income is higher than they expected). In a parallel but unrelated development Maccheroni, Marinacci, and Rustichini (2009a,b) have also proposed that the utility of a target amount depends upon a comparison with a reference distribution. Here, the reference distribution is the set of outcomes experienced by others. An agent's utility is a combination of classical, subjective expected utility and a position index which increases with the utility of the agent's own outcome and decreases with the distribution of the agents' peers.

These models are related to the decision-by-sampling model (Stewart, Chater, and Brown, 2006) described earlier, in which the subjective value of an amount (or probability or delay) is just the rank position within the reference distribution. Both Kőszegi and Rabin and Maccheroni, Marinacci, and Rustichini assume an additional, stable classical utility component. And Kőszegi and Rabin assume that their gain-loss utility is cardinal, rather than just the rank position in the distribution. The models also differ on the source of the reference distribution: Kőszegi and Rabin assume the distribution is of (rational) expected outcomes, Maccheroni, Marinacci, and Rustichini assume it is the outcomes of peers, and Stewart, Chater, and Brown assume it is the attribute values in immediate memory, many of which are from the immediate context or are evoked from long-term memory by the immediate context. This paper, and related work, provide experimental and field evidence that it is the attribute values in mind that matter as decision by sampling states—and these values could well include expected outcomes or the outcomes of others.

7.3.2. *Ordinal utility*

Read (2007, see also Stigler, 1950a,b) gives a review of the historical shift from Bentham's cardinal utility to the modern *ordinal* utility approach. The problem with Bentham's approach is that the utility of a good is not independent of other goods as is assumed in Bentham's approach. For example, coffee is more useful if one possesses an espresso machine. Edgworth proposed that utility should be a function of the entire basket of goods one possesses. Pareto demonstrated that all that is required for modern economics is that people can order their preferences over baskets of goods. Baskets of goods can be represented in a multidimensional space, with baskets between which the agent is indifferent joined by contours of equal utility. Now utility is ordinal: All that is required for modern economics is that, if one basket is preferred to another, then it comes from a higher indifference curve. Preferences are said to be *revealed*. Numbers attached to the contours are only ordinal and are arbitrary as long as higher indifference curves are assigned higher numbers. The difference between utilities has no meaning, and thus any monotonic transformation of a utility function is permitted. Given our experimental manipulations do not change the ordering of utilities—utility curves always increase monotonically—one might argue that these data do not present a problem for the modern economic approach.

But we believe our results are still important for modern economics. Essentially, risky choice can be viewed as choice between baskets of risk and reward, where one basket might contain a large reward and a large risk whereas another might contain a smaller reward at smaller risk. Intertemporal choice can be viewed in the same way. Our results suggest that it should be possible to observe a reversal of the preference between two "baskets" in difference contexts. This is just the result found by Stewart (2009) and Ungemach, Stewart, and Reimers (2011) described earlier (see also Dorlet, Simonson, and Tversky, 2000). Thus the notion that people have stable preferences between baskets of goods is, empirically, not true.

7.3.3. *Neuroeconomic evidence*

Our experimental results are consistent with recent evidence of coding of relative value within the brain. Seymour and McClure (2008) give a recent review. For example, Tremblay and Schultz (1999) recorded from single cells in the macaque orbitofrontal cortex. Cells fired more strongly on presentation of a piece of apple as a reward when the other available reward was a piece of cereal (which monkeys do not like that much) compared to when the other available reward was a raisin (which monkeys love): The value of the apple was coded relative to the other available rewards.

Some studies in neuroeconomics seem to provide compelling evidence for utility or value functions. For example, Tom, Fox, Trepel, and Poldrack (2007) showed that regions of the striatum and VMPFC show conjoint positive sensitivity to increases in the size of a gain and decreases in the size of a loss. Further, the majority of individuals showed loss aversion: the increase in activity for a given gain was about half the decrease in activity for the same sized loss. Tom, Fox, Trepel, and Poldrack argue that they have uncovered the neural basis of the value function and loss aversion. But this evidence is completely consistent with our rank-based explanation of the present data. Tom, Fox, Trepel, and Poldrack used gains in the range \$10 to \$40, but losses in the range \$5 to \$20 in their experiment. Thus finding that people are twice as sensitive to losses is entirely expected given the distribution of losses spans only half the range that the distribution of gains covers. What is required to demonstrate the existence of stable value functions would be to show that the brain activation for a particular gain (or loss) is invariant across a set of contexts. To our knowledge, such an experiment has not been done.

7.4. *Conclusion*

To close, if we can measure lots of utility functions, which utility function is the right one? Alternatively, if one infers a level of risk aversion from a series of choices, whether in an experimental setting or by observing real-world decision, which choices reveal the “true” level of risk aversion? One can pose similar questions for weighting and discounting functions. The answer, we fear, is that there is no method which gives, even with careful counterbalancing, the true level of risk aversion or the true shape of a utility function. In any given situation, one can observe choices and infer a shape or level of risk aversion. But as soon as the context changes—that is, as soon as the decision maker experiences any new amount—the measured shape or level of risk aversion will no longer apply. Similarly, one cannot examine risk aversion in one market situation and generalize to another. Bluntly: individuals do not have utility functions with particular shapes and thus particular levels of risk aversion. If it is meaningful to talk of utility or risk aversion at all, it is only in the instantaneous interaction between the distribution of attribute values in the immediate context and the decision maker.

Yet laboratory experiments and observations from the field consistently suggest concave utility functions, inverse-S-shaped probability weighting functions, and hyperbolic-like discounting functions. Here we have offered the decision by sampling theory as an explanation of these findings: The interaction between the distribtion of gains, losses, risks, and delays in the world (or the laboratory) with the simple cognitive tools of sampling from memory, binary ordinal comparison, and frequency accumulation gives the ubiquitous psychoeconomic functions. We have presented four laboratory experiments testing this account: Systematic manipulation of the distributions of gains, risks, and delays gives large and systematic change in the utility, weighting, and discounting functions revealed from people’s choices. Decision by sampling provides a common account of the origin of utility, weighting, and discounting functions—and the explanation of how we, as architects of the environment in which decisions were made, were able to change their shapes.

APPENDIX A: CHOICES

Table A.1: Experiment 1 Choices

	Positive Skew				Negative Skew			
ID	p	x	q	y	p	x	q	y
1	20%	£20	40%	£10	20%	£310	40%	£10
2	20%	£20	60%	£10	20%	£310	60%	£10
3	20%	£20	80%	£10	20%	£310	80%	£10
4	20%	£20	100%	£10	20%	£310	100%	£10
5	40%	£20	60%	£10	40%	£310	60%	£10
6	40%	£20	80%	£10	40%	£310	80%	£10
7	40%	£20	100%	£10	40%	£310	100%	£10
8	60%	£20	80%	£10	60%	£310	80%	£10
9	60%	£20	100%	£10	60%	£310	100%	£10
10	80%	£20	100%	£10	80%	£310	100%	£10
11	20%	£50	40%	£10	20%	£410	40%	£10
12	20%	£50	60%	£10	20%	£410	60%	£10
13	20%	£50	80%	£10	20%	£410	80%	£10
14	20%	£50	100%	£10	20%	£410	100%	£10
15	20%	£50	40%	£20	20%	£410	40%	£310
16	20%	£50	60%	£20	20%	£410	60%	£310
17	20%	£50	80%	£20	20%	£410	80%	£310
18	20%	£50	100%	£20	20%	£410	100%	£310
19	40%	£50	60%	£10	40%	£410	60%	£10
20	40%	£50	80%	£10	40%	£410	80%	£10
21	40%	£50	100%	£10	40%	£410	100%	£10
22	40%	£50	60%	£20	40%	£410	60%	£310
23	40%	£50	80%	£20	40%	£410	80%	£310
24	40%	£50	100%	£20	40%	£410	100%	£310
25	60%	£50	80%	£10	60%	£410	80%	£10
26	60%	£50	100%	£10	60%	£410	100%	£10
27	60%	£50	80%	£20	60%	£410	80%	£310
28	60%	£50	100%	£20	60%	£410	100%	£310
29	80%	£50	100%	£10	80%	£410	100%	£10
30	80%	£50	100%	£20	80%	£410	100%	£310
31	20%	£100	40%	£10	20%	£460	40%	£10
32	20%	£100	60%	£10	20%	£460	60%	£10
33	20%	£100	80%	£10	20%	£460	80%	£10
34	20%	£100	100%	£10	20%	£460	100%	£10
35	20%	£100	40%	£20	20%	£460	40%	£310
36	20%	£100	60%	£20	20%	£460	60%	£310
37	20%	£100	80%	£20	20%	£460	80%	£310
38	20%	£100	100%	£20	20%	£460	100%	£310

continued on next page

Table A.1 (continued from previous page)

ID	Positive Skew *p*	*x*	*q*	*y*	Negative Skew *p*	*x*	*q*	*y*
39	20%	£100	40%	£50	20%	£460	40%	£410
40	20%	£100	60%	£50	20%	£460	60%	£410
41	20%	£100	80%	£50	20%	£460	80%	£410
42	20%	£100	100%	£50	20%	£460	100%	£410
43	40%	£100	60%	£10	40%	£460	60%	£10
44	40%	£100	80%	£10	40%	£460	80%	£10
45	40%	£100	100%	£10	40%	£460	100%	£10
46	40%	£100	60%	£20	40%	£460	60%	£310
47	40%	£100	80%	£20	40%	£460	80%	£310
48	40%	£100	100%	£20	40%	£460	100%	£310
49	40%	£100	60%	£50	40%	£460	60%	£410
50	40%	£100	80%	£50	40%	£460	80%	£410
51	40%	£100	100%	£50	40%	£460	100%	£410
52	60%	£100	80%	£10	60%	£460	80%	£10
53	60%	£100	100%	£10	60%	£460	100%	£10
54	60%	£100	80%	£20	60%	£460	80%	£310
55	60%	£100	100%	£20	60%	£460	100%	£310
56	60%	£100	80%	£50	60%	£460	80%	£410
57	60%	£100	100%	£50	60%	£460	100%	£410
58	80%	£100	100%	£10	80%	£460	100%	£10
59	80%	£100	100%	£20	80%	£460	100%	£310
60	80%	£100	100%	£50	80%	£460	100%	£410
61	20%	£200	40%	£10	20%	£490	40%	£10
62	20%	£200	60%	£10	20%	£490	60%	£10
63	20%	£200	80%	£10	20%	£490	80%	£10
64	20%	£200	100%	£10	20%	£490	100%	£10
65	20%	£200	40%	£20	20%	£490	40%	£310
66	20%	£200	60%	£20	20%	£490	60%	£310
67	20%	£200	80%	£20	20%	£490	80%	£310
68	20%	£200	100%	£20	20%	£490	100%	£310
69	20%	£200	40%	£50	20%	£490	40%	£410
70	20%	£200	60%	£50	20%	£490	60%	£410
71	20%	£200	80%	£50	20%	£490	80%	£410
72	20%	£200	100%	£50	20%	£490	100%	£410
73	20%	£200	40%	£100	20%	£490	40%	£460
74	20%	£200	60%	£100	20%	£490	60%	£460
75	20%	£200	80%	£100	20%	£490	80%	£460
76	20%	£200	100%	£100	20%	£490	100%	£460
77	40%	£200	60%	£10	40%	£490	60%	£10

continued on next page

Table A.1 (continued from previous page)

ID	Positive Skew p	x	q	y	Negative Skew p	x	q	y
78	40%	£200	80%	£10	40%	£490	80%	£10
79	40%	£200	100%	£10	40%	£490	100%	£10
80	40%	£200	60%	£20	40%	£490	60%	£310
81	40%	£200	80%	£20	40%	£490	80%	£310
82	40%	£200	100%	£20	40%	£490	100%	£310
83	40%	£200	60%	£50	40%	£490	60%	£410
84	40%	£200	80%	£50	40%	£490	80%	£410
85	40%	£200	100%	£50	40%	£490	100%	£410
86	40%	£200	60%	£100	40%	£490	60%	£460
87	40%	£200	80%	£100	40%	£490	80%	£460
88	40%	£200	100%	£100	40%	£490	100%	£460
89	60%	£200	80%	£10	60%	£490	80%	£10
90	60%	£200	100%	£10	60%	£490	100%	£10
91	60%	£200	80%	£20	60%	£490	80%	£310
92	60%	£200	100%	£20	60%	£490	100%	£310
93	60%	£200	80%	£50	60%	£490	80%	£410
94	60%	£200	100%	£50	60%	£490	100%	£410
95	60%	£200	80%	£100	60%	£490	80%	£460
96	60%	£200	100%	£100	60%	£490	100%	£460
97	80%	£200	100%	£10	80%	£490	100%	£10
98	80%	£200	100%	£20	80%	£490	100%	£310
99	80%	£200	100%	£50	80%	£490	100%	£410
100	80%	£200	100%	£100	80%	£490	100%	£460
101	20%	£500	40%	£10	20%	£500	40%	£10
102	20%	£500	60%	£10	20%	£500	60%	£10
103	20%	£500	80%	£10	20%	£500	80%	£10
104	20%	£500	100%	£10	20%	£500	100%	£10
105	20%	£500	40%	£20	20%	£500	40%	£310
106	20%	£500	60%	£20	20%	£500	60%	£310
107	20%	£500	80%	£20	20%	£500	80%	£310
108	20%	£500	100%	£20	20%	£500	100%	£310
109	20%	£500	40%	£50	20%	£500	40%	£410
110	20%	£500	60%	£50	20%	£500	60%	£410
111	20%	£500	80%	£50	20%	£500	80%	£410
112	20%	£500	100%	£50	20%	£500	100%	£410
113	20%	£500	40%	£100	20%	£500	40%	£460
114	20%	£500	60%	£100	20%	£500	60%	£460
115	20%	£500	80%	£100	20%	£500	80%	£460
116	20%	£500	100%	£100	20%	£500	100%	£460

continued on next page

Table A.1 (continued from previous page)

	Positive Skew				Negative Skew			
ID	*p*	*x*	*q*	*y*	*p*	*x*	*q*	*y*
117	20%	£500	40%	£200	20%	£500	40%	£490
118	20%	£500	60%	£200	20%	£500	60%	£490
119	20%	£500	80%	£200	20%	£500	80%	£490
120	20%	£500	100%	£200	20%	£500	100%	£490
121	40%	£500	60%	£10	40%	£500	60%	£10
122	40%	£500	80%	£10	40%	£500	80%	£10
123	40%	£500	100%	£10	40%	£500	100%	£10
124	40%	£500	60%	£20	40%	£500	60%	£310
125	40%	£500	80%	£20	40%	£500	80%	£310
126	40%	£500	100%	£20	40%	£500	100%	£310
127	40%	£500	60%	£50	40%	£500	60%	£410
128	40%	£500	80%	£50	40%	£500	80%	£410
129	40%	£500	100%	£50	40%	£500	100%	£410
130	40%	£500	60%	£100	40%	£500	60%	£460
131	40%	£500	80%	£100	40%	£500	80%	£460
132	40%	£500	100%	£100	40%	£500	100%	£460
133	40%	£500	60%	£200	40%	£500	60%	£490
134	40%	£500	80%	£200	40%	£500	80%	£490
135	40%	£500	100%	£200	40%	£500	100%	£490
136	60%	£500	80%	£10	60%	£500	80%	£10
137	60%	£500	100%	£10	60%	£500	100%	£10
138	60%	£500	80%	£20	60%	£500	80%	£310
139	60%	£500	100%	£20	60%	£500	100%	£310
140	60%	£500	80%	£50	60%	£500	80%	£410
141	60%	£500	100%	£50	60%	£500	100%	£410
142	60%	£500	80%	£100	60%	£500	80%	£460
143	60%	£500	100%	£100	60%	£500	100%	£460
144	60%	£500	80%	£200	60%	£500	80%	£490
145	60%	£500	100%	£200	60%	£500	100%	£490
146	80%	£500	100%	£10	80%	£500	100%	£10
147	80%	£500	100%	£20	80%	£500	100%	£310
148	80%	£500	100%	£50	80%	£500	100%	£410
149	80%	£500	100%	£100	80%	£500	100%	£460
150	80%	£500	100%	£200	80%	£500	100%	£490
151	100%	£500	80%	£20	100%	£500	80%	£310
152	100%	£100	100%	£20	100%	£460	100%	£310
153	100%	£10	20%	£10	100%	£10	20%	£10
154	80%	£500	40%	£20	80%	£500	40%	£310
155	20%	£500	20%	£100	20%	£500	20%	£460

continued on next page

Table A.1 (continued from previous page)

	Positive Skew				Negative Skew			
ID	*p*	*x*	*q*	*y*	*p*	*x*	*q*	*y*
156	100%	£100	20%	£100	100%	£460	20%	£460
157	80%	£200	80%	£20	80%	£490	80%	£310
158	40%	£100	20%	£10	40%	£460	20%	£10
159	100%	£200	20%	£100	100%	£490	20%	£460
160	60%	£500	20%	£500	60%	£500	20%	£500
161	80%	£200	40%	£100	80%	£490	40%	£460
162	80%	£200	80%	£100	80%	£490	80%	£460
163	40%	£200	40%	£10	40%	£490	40%	£10
164	100%	£200	100%	£100	100%	£490	100%	£460
165	20%	£50	20%	£10	20%	£410	20%	£10
166	60%	£200	20%	£10	60%	£490	20%	£10
167	100%	£200	100%	£10	100%	£490	100%	£10
168	100%	£10	80%	£10	100%	£10	80%	£10
169	20%	£500	20%	£50	20%	£500	20%	£410
170	80%	£500	40%	£500	80%	£500	40%	£500
171	60%	£20	20%	£10	60%	£310	20%	£10
172	80%	£20	60%	£10	80%	£310	60%	£10
173	40%	£200	20%	£50	40%	£490	20%	£410
174	80%	£500	60%	£20	80%	£500	60%	£310
175	100%	£200	40%	£50	100%	£490	40%	£410
176	60%	£20	20%	£20	60%	£310	20%	£310
177	100%	£200	80%	£50	100%	£490	80%	£410
178	60%	£500	60%	£200	60%	£500	60%	£490
179	100%	£200	20%	£20	100%	£490	20%	£310
180	60%	£50	20%	£50	60%	£410	20%	£410

Table A.2: Experiment 2 Choices

	Positive Skew				Negative Skew			
ID	*p*	*x*	*q*	*y*	*p*	*x*	*q*	*y*
1	20%	£200	40%	£100	30%	£200	70%	£100
2	40%	£400	70%	£100	70%	£400	80%	£100
3	10%	£500	40%	£200	10%	£500	70%	£200
4	70%	£300	90%	£100	80%	£300	90%	£100
5	30%	£300	70%	£100	60%	£300	80%	£100
6	10%	£400	90%	£300	10%	£400	90%	£300
7	10%	£400	40%	£100	10%	£400	70%	£100

continued on next page

Table A.2 (continued from previous page)

ID	Positive Skew *p*	*x*	*q*	*y*	Negative Skew *p*	*x*	*q*	*y*
8	30%	£400	40%	£300	60%	£400	70%	£300
9	10%	£400	20%	£100	10%	£400	30%	£100
10	20%	£500	70%	£400	30%	£500	80%	£400
11	40%	£500	70%	£400	70%	£500	80%	£400
12	30%	£400	70%	£300	60%	£400	80%	£300
13	30%	£300	40%	£100	60%	£300	70%	£100
14	20%	£400	30%	£200	30%	£400	60%	£200
15	20%	£500	30%	£400	30%	£500	60%	£400
16	10%	£300	70%	£200	10%	£300	80%	£200
17	30%	£400	70%	£200	60%	£400	80%	£200
18	10%	£500	30%	£200	10%	£500	60%	£200
19	20%	£400	40%	£100	30%	£400	70%	£100
20	20%	£400	40%	£200	30%	£400	70%	£200
21	10%	£300	40%	£100	10%	£300	70%	£100
22	10%	£400	70%	£300	10%	£400	80%	£300
23	30%	£500	70%	£200	60%	£500	80%	£200
24	40%	£300	90%	£100	70%	£300	90%	£100
25	20%	£500	90%	£400	30%	£500	90%	£400
26	10%	£400	20%	£200	10%	£400	30%	£200
27	10%	£300	70%	£100	10%	£300	80%	£100
28	40%	£200	90%	£100	70%	£200	90%	£100
29	70%	£400	90%	£200	80%	£400	90%	£200
30	10%	£200	90%	£100	10%	£200	90%	£100
31	20%	£400	70%	£200	30%	£400	80%	£200
32	20%	£500	40%	£100	30%	£500	70%	£100
33	10%	£500	20%	£100	10%	£500	30%	£100
34	70%	£400	90%	£100	80%	£400	90%	£100
35	30%	£500	90%	£200	60%	£500	90%	£200
36	40%	£400	70%	£300	70%	£400	80%	£300
37	40%	£400	90%	£200	70%	£400	90%	£200
38	20%	£400	30%	£300	30%	£400	60%	£300
39	70%	£200	90%	£100	80%	£200	90%	£100
40	30%	£300	40%	£200	60%	£300	70%	£200
41	10%	£500	70%	£400	10%	£500	80%	£400
42	10%	£300	90%	£100	10%	£300	90%	£100
43	20%	£300	90%	£100	30%	£300	90%	£100
44	20%	£500	40%	£400	30%	£500	70%	£400
45	40%	£400	90%	£100	70%	£400	90%	£100
46	40%	£300	70%	£200	70%	£300	80%	£200

continued on next page

Table A.2 (continued from previous page)

ID	Positive Skew p	x	q	y	Negative Skew p	x	q	y
47	20%	£500	70%	£200	30%	£500	80%	£200
48	30%	£200	90%	£100	60%	£200	90%	£100
49	30%	£300	70%	£200	60%	£300	80%	£200
50	10%	£500	70%	£300	10%	£500	80%	£300
51	20%	£400	90%	£200	30%	£400	90%	£200
52	40%	£500	70%	£300	70%	£500	80%	£300
53	30%	£500	70%	£300	60%	£500	80%	£300
54	10%	£500	90%	£300	10%	£500	90%	£300
55	30%	£200	70%	£100	60%	£200	80%	£100
56	20%	£500	90%	£200	30%	£500	90%	£200
57	10%	£400	30%	£200	10%	£400	60%	£200
58	40%	£500	90%	£300	70%	£500	90%	£300
59	10%	£500	90%	£100	10%	£500	90%	£100
60	20%	£200	90%	£100	30%	£200	90%	£100
61	40%	£500	70%	£200	70%	£500	80%	£200
62	40%	£500	90%	£400	70%	£500	90%	£400
63	40%	£300	70%	£100	70%	£300	80%	£100
64	10%	£200	20%	£100	10%	£200	30%	£100
65	20%	£300	90%	£200	30%	£300	90%	£200
66	10%	£200	40%	£100	10%	£200	70%	£100
67	10%	£500	90%	£200	10%	£500	90%	£200
68	20%	£500	70%	£100	30%	£500	80%	£100
69	20%	£400	70%	£100	30%	£400	80%	£100
70	30%	£500	90%	£100	60%	£500	90%	£100
71	40%	£400	70%	£200	70%	£400	80%	£200
72	20%	£500	90%	£300	30%	£500	90%	£300
73	30%	£400	40%	£200	60%	£400	70%	£200
74	30%	£500	40%	£100	60%	£500	70%	£100
75	30%	£400	70%	£100	60%	£400	80%	£100
76	20%	£500	30%	£200	30%	£500	60%	£200
77	40%	£200	70%	£100	70%	£200	80%	£100
78	10%	£300	90%	£200	10%	£300	90%	£200
79	10%	£300	30%	£100	10%	£300	60%	£100
80	10%	£300	30%	£200	10%	£300	60%	£200
81	10%	£400	40%	£200	10%	£400	70%	£200
82	30%	£300	90%	£100	60%	£300	90%	£100
83	20%	£500	40%	£300	30%	£500	70%	£300
84	30%	£400	90%	£100	60%	£400	90%	£100
85	10%	£500	70%	£100	10%	£500	80%	£100

continued on next page

Table A.2 (continued from previous page)

	Positive Skew				Negative Skew			
ID	p	x	q	y	p	x	q	y
86	40%	£400	90%	£300	70%	£400	90%	£300
87	30%	£300	90%	£200	60%	£300	90%	£200
88	20%	£500	90%	£100	30%	£500	90%	£100
89	20%	£400	30%	£100	30%	£400	60%	£100
90	30%	£400	90%	£200	60%	£400	90%	£200
91	20%	£400	40%	£300	30%	£400	70%	£300
92	20%	£400	90%	£300	30%	£400	90%	£300
93	20%	£400	70%	£300	30%	£400	80%	£300
94	30%	£500	90%	£400	60%	£500	90%	£400
95	70%	£300	90%	£200	80%	£300	90%	£200
96	20%	£500	40%	£200	30%	£500	70%	£200
97	20%	£300	40%	£100	30%	£300	70%	£100
98	20%	£300	40%	£200	30%	£300	70%	£200
99	70%	£500	90%	£400	80%	£500	90%	£400
100	10%	£500	20%	£400	10%	£500	30%	£400
101	10%	£400	90%	£200	10%	£400	90%	£200
102	30%	£500	40%	£400	60%	£500	70%	£400
103	10%	£500	40%	£400	10%	£500	70%	£400
104	70%	£500	90%	£300	80%	£500	90%	£300
105	10%	£500	90%	£400	10%	£500	90%	£400
106	30%	£500	40%	£300	60%	£500	70%	£300
107	70%	£500	90%	£100	80%	£500	90%	£100
108	70%	£500	90%	£200	80%	£500	90%	£200
109	20%	£300	70%	£100	30%	£300	80%	£100
110	10%	£300	20%	£100	10%	£300	30%	£100
111	10%	£500	40%	£100	10%	£500	70%	£100
112	10%	£200	30%	£100	10%	£200	60%	£100
113	10%	£300	20%	£200	10%	£300	30%	£200
114	10%	£500	70%	£200	10%	£500	80%	£200
115	20%	£400	90%	£100	30%	£400	90%	£100
116	10%	£400	70%	£100	10%	£400	80%	£100
117	20%	£300	30%	£200	30%	£300	60%	£200
118	30%	£500	40%	£200	60%	£500	70%	£200
119	10%	£500	40%	£300	10%	£500	70%	£300
120	40%	£500	90%	£100	70%	£500	90%	£100
121	20%	£500	10%	£400	30%	£500	10%	£400
122	90%	£200	70%	£200	90%	£200	80%	£200
123	70%	£500	10%	£300	80%	£500	10%	£300
124	40%	£400	30%	£300	70%	£400	60%	£300

continued on next page

Table A.2 (continued from previous page)

ID	Positive Skew p	x	q	y	Negative Skew p	x	q	y
125	40%	£500	30%	£200	70%	£500	60%	£200
126	90%	£300	90%	£200	90%	£300	90%	£200
127	90%	£400	20%	£300	90%	£400	30%	£300
128	90%	£500	30%	£300	90%	£500	60%	£300
129	90%	£500	70%	£400	90%	£500	80%	£400
130	90%	£300	70%	£100	90%	£300	80%	£100
131	10%	£400	10%	£200	10%	£400	10%	£200
132	30%	£500	20%	£100	60%	£500	30%	£100
133	30%	£500	30%	£400	60%	£500	60%	£400
134	20%	£100	10%	£100	30%	£100	10%	£100
135	30%	£500	10%	£400	60%	£500	10%	£400
136	20%	£200	10%	£200	30%	£200	10%	£200
137	20%	£400	10%	£200	30%	£400	10%	£200
138	40%	£400	10%	£300	70%	£400	10%	£300
139	70%	£200	30%	£100	80%	£200	60%	£100
140	30%	£500	20%	£300	60%	£500	30%	£300
141	90%	£500	30%	£500	90%	£500	60%	£500
142	70%	£300	70%	£200	80%	£300	80%	£200
143	40%	£300	20%	£200	70%	£300	30%	£200
144	30%	£300	20%	£100	60%	£300	30%	£100
145	90%	£500	40%	£100	90%	£500	70%	£100
146	90%	£400	70%	£200	90%	£400	80%	£200
147	90%	£500	10%	£400	90%	£500	10%	£400
148	70%	£400	40%	£200	80%	£400	70%	£200
149	70%	£300	10%	£200	80%	£300	10%	£200
150	90%	£500	90%	£100	90%	£500	90%	£100

Table A.3: Experiment 3 Choices

ID	Positive Skew x_1	t_1	x_2	t_2	Negative Skew x_1	t_1	x_2	t_2
1	£500	2 months	£20	2 days	£500	2 months	£310	2 days
2	£500	1 month	£50	2 weeks	£500	1 month	£410	2 weeks
3	£500	2 months	£50	1 day	£500	2 months	£410	1 day
4	£200	1 year	£10	2 weeks	£490	1 year	£10	2 weeks
5	£50	6 months	£20	2 weeks	£410	6 months	£310	2 weeks
6	£100	2 months	£50	1 day	£460	2 months	£410	1 day

continued on next page

Table A.3 (continued from previous page)

	Positive Skew				Negative Skew			
ID	x_1	t_1	x_2	t_2	x_1	t_1	x_2	t_2
7	£100	6 months	£10	1 week	£460	6 months	£10	1 week
8	£500	1 month	£200	1 day	£500	1 month	£490	1 day
9	£500	6 months	£10	2 weeks	£500	6 months	£10	2 weeks
10	£500	2 months	£200	2 weeks	£500	2 months	£490	2 weeks
11	£500	2 months	£20	1 week	£500	2 months	£310	1 week
12	£100	1 year	£10	2 weeks	£460	1 year	£10	2 weeks
13	£50	2 months	£10	2 days	£410	2 months	£10	2 days
14	£200	6 months	£100	2 months	£490	6 months	£460	2 months
15	£20	6 months	£10	2 months	£310	6 months	£10	2 months
16	£200	2 weeks	£100	1 day	£490	2 weeks	£460	1 day
17	£200	1 month	£10	1 week	£490	1 month	£10	1 week
18	£500	2 months	£50	1 month	£500	2 months	£410	1 month
19	£200	2 months	£10	2 days	£490	2 months	£10	2 days
20	£200	2 months	£20	2 weeks	£490	2 months	£310	2 weeks
21	£500	1 week	£200	2 days	£500	1 week	£490	2 days
22	£500	1 month	£20	1 day	£500	1 month	£310	1 day
23	£100	6 months	£10	1 month	£460	6 months	£10	1 month
24	£200	2 weeks	£50	1 week	£490	2 weeks	£410	1 week
25	£200	2 days	£10	1 day	£490	2 days	£10	1 day
26	£500	2 months	£200	1 day	£500	2 months	£490	1 day
27	£20	1 month	£10	1 week	£310	1 month	£10	1 week
28	£100	6 months	£10	2 months	£460	6 months	£10	2 months
29	£100	2 days	£50	1 day	£460	2 days	£410	1 day
30	£50	2 months	£20	2 weeks	£410	2 months	£310	2 weeks
31	£100	2 weeks	£20	2 days	£460	2 weeks	£310	2 days
32	£200	2 months	£100	1 month	£490	2 months	£460	1 month
33	£500	6 months	£20	2 months	£500	6 months	£310	2 months
34	£200	2 months	£50	1 day	£490	2 months	£410	1 day
35	£50	6 months	£20	1 day	£410	6 months	£310	1 day
36	£500	1 month	£100	2 days	£500	1 month	£460	2 days
37	£100	1 year	£20	1 week	£460	1 year	£310	1 week
38	£500	6 months	£200	2 weeks	£500	6 months	£490	2 weeks
39	£500	6 months	£50	2 weeks	£500	6 months	£410	2 weeks
40	£200	1 month	£10	2 weeks	£490	1 month	£10	2 weeks
41	£500	1 year	£10	2 months	£500	1 year	£10	2 months
42	£20	2 months	£10	1 month	£310	2 months	£10	1 month
43	£20	2 months	£10	2 weeks	£310	2 months	£10	2 weeks
44	£500	6 months	£200	2 months	£500	6 months	£490	2 months
45	£200	1 year	£10	1 month	£490	1 year	£10	1 month

continued on next page

Table A.3 (continued from previous page)

ID	Positive Skew x_1	t_1	x_2	t_2	Negative Skew x_1	t_1	x_2	t_2
46	£500	1 year	£20	1 week	£500	1 year	£310	1 week
47	£200	6 months	£100	2 days	£490	6 months	£460	2 days
48	£200	2 months	£10	1 day	£490	2 months	£10	1 day
49	£100	2 weeks	£50	2 days	£460	2 weeks	£410	2 days
50	£200	6 months	£100	1 day	£490	6 months	£460	1 day
51	£20	1 month	£10	2 days	£310	1 month	£10	2 days
52	£100	2 weeks	£20	1 week	£460	2 weeks	£310	1 week
53	£200	1 month	£100	1 day	£490	1 month	£460	1 day
54	£100	6 months	£20	1 month	£460	6 months	£310	1 month
55	£200	2 days	£50	1 day	£490	2 days	£410	1 day
56	£500	6 months	£50	2 days	£500	6 months	£410	2 days
57	£50	1 year	£20	2 weeks	£410	1 year	£310	2 weeks
58	£20	1 year	£10	2 days	£310	1 year	£10	2 days
59	£200	2 weeks	£50	1 day	£490	2 weeks	£410	1 day
60	£500	1 year	£100	1 week	£500	1 year	£460	1 week
61	£500	2 months	£10	1 day	£500	2 months	£10	1 day
62	£20	1 year	£10	1 month	£310	1 year	£10	1 month
63	£200	2 months	£50	2 weeks	£490	2 months	£410	2 weeks
64	£100	1 year	£10	2 days	£460	1 year	£10	2 days
65	£200	6 months	£50	1 week	£490	6 months	£410	1 week
66	£100	1 year	£20	1 day	£460	1 year	£310	1 day
67	£500	2 weeks	£50	2 days	£500	2 weeks	£410	2 days
68	£200	1 month	£20	1 week	£490	1 month	£310	1 week
69	£50	2 days	£10	1 day	£410	2 days	£10	1 day
70	£100	2 weeks	£50	1 week	£460	2 weeks	£410	1 week
71	£500	6 months	£10	1 day	£500	6 months	£10	1 day
72	£50	6 months	£10	2 months	£410	6 months	£10	2 months
73	£500	1 year	£20	1 month	£500	1 year	£310	1 month
74	£50	1 month	£20	2 weeks	£410	1 month	£310	2 weeks
75	£200	2 weeks	£50	2 days	£490	2 weeks	£410	2 days
76	£200	1 week	£100	1 day	£490	1 week	£460	1 day
77	£500	6 months	£100	2 months	£500	6 months	£460	2 months
78	£500	1 week	£50	1 day	£500	1 week	£410	1 day
79	£500	6 months	£50	1 day	£500	6 months	£410	1 day
80	£500	2 weeks	£200	1 week	£500	2 weeks	£490	1 week
81	£200	6 months	£50	2 weeks	£490	6 months	£410	2 weeks
82	£200	1 week	£10	1 day	£490	1 week	£10	1 day
83	£100	2 months	£10	1 week	£460	2 months	£10	1 week
84	£50	1 week	£20	2 days	£410	1 week	£310	2 days

continued on next page

Table A.3 (continued from previous page)

ID	Positive Skew x_1	t_1	x_2	t_2	Negative Skew x_1	t_1	x_2	t_2
85	£100	6 months	£10	1 day	£460	6 months	£10	1 day
86	£50	1 week	£10	1 day	£410	1 week	£10	1 day
87	£500	2 weeks	£50	1 week	£500	2 weeks	£410	1 week
88	£20	6 months	£10	1 day	£310	6 months	£10	1 day
89	£500	1 week	£100	2 days	£500	1 week	£460	2 days
90	£100	1 year	£50	6 months	£460	1 year	£410	6 months
91	£20	2 weeks	£10	1 day	£310	2 weeks	£10	1 day
92	£200	2 months	£100	2 weeks	£490	2 months	£460	2 weeks
93	£20	1 year	£10	2 weeks	£310	1 year	£10	2 weeks
94	£100	1 year	£50	1 week	£460	1 year	£410	1 week
95	£100	1 year	£10	1 week	£460	1 year	£10	1 week
96	£500	2 weeks	£100	2 days	£500	2 weeks	£460	2 days
97	£500	2 months	£100	2 days	£500	2 months	£460	2 days
98	£100	2 months	£20	2 days	£460	2 months	£310	2 days
99	£100	1 week	£10	1 day	£460	1 week	£10	1 day
100	£500	6 months	£100	1 week	£500	6 months	£460	1 week
101	£50	1 year	£10	1 week	£410	1 year	£10	1 week
102	£500	1 month	£200	2 days	£500	1 month	£490	2 days
103	£50	2 months	£10	2 weeks	£410	2 months	£10	2 weeks
104	£500	1 year	£50	1 day	£500	1 year	£410	1 day
105	£50	1 year	£20	6 months	£410	1 year	£310	6 months
106	£200	1 week	£20	2 days	£490	1 week	£310	2 days
107	£500	1 year	£200	1 day	£500	1 year	£490	1 day
108	£500	6 months	£100	1 month	£500	6 months	£460	1 month
109	£200	1 year	£100	2 weeks	£490	1 year	£460	2 weeks
110	£500	2 months	£200	1 month	£500	2 months	£490	1 month
111	£200	6 months	£50	1 month	£490	6 months	£410	1 month
112	£500	6 months	£200	1 week	£500	6 months	£490	1 week
113	£100	2 months	£10	2 weeks	£460	2 months	£10	2 weeks
114	£200	1 month	£50	2 weeks	£490	1 month	£410	2 weeks
115	£500	1 week	£10	1 day	£500	1 week	£10	1 day
116	£500	1 year	£100	6 months	£500	1 year	£460	6 months
117	£200	2 months	£10	1 month	£490	2 months	£10	1 month
118	£500	2 months	£10	1 month	£500	2 months	£10	1 month
119	£500	2 months	£200	1 week	£500	2 months	£490	1 week
120	£50	6 months	£20	1 week	£410	6 months	£310	1 week
121	£200	1 month	£200	2 weeks	£490	1 month	£490	2 weeks
122	£50	1 month	£20	2 months	£410	1 month	£310	2 months
123	£100	2 days	£10	2 weeks	£460	2 days	£10	2 weeks

continued on next page

Table A.3 (continued from previous page)

	Positive Skew				Negative Skew			
ID	x_1	t_1	x_2	t_2	x_1	t_1	x_2	t_2
124	£100	6 months	£10	1 year	£460	6 months	£10	1 year
125	£500	1 day	£200	1 week	£500	1 day	£490	1 week
126	£500	2 months	£100	1 year	£500	2 months	£460	1 year
127	£500	2 days	£10	1 week	£500	2 days	£10	1 week
128	£500	1 year	£500	1 month	£500	1 year	£500	1 month
129	£20	2 days	£20	1 day	£310	2 days	£310	1 day
130	£100	2 months	£10	6 months	£460	2 months	£10	6 months
131	£200	1 day	£20	2 weeks	£490	1 day	£310	2 weeks
132	£10	6 months	£10	2 months	£10	6 months	£10	2 months
133	£20	1 year	£20	1 month	£310	1 year	£310	1 month
134	£500	1 week	£200	2 weeks	£500	1 week	£490	2 weeks
135	£10	2 months	£10	1 month	£10	2 months	£10	1 month
136	£100	2 weeks	£20	2 weeks	£460	2 weeks	£310	2 weeks
137	£200	1 week	£10	6 months	£490	1 week	£10	6 months
138	£500	2 days	£200	6 months	£500	2 days	£490	6 months
139	£100	1 day	£10	1 week	£460	1 day	£10	1 week
140	£200	6 months	£200	2 days	£490	6 months	£490	2 days
141	£50	2 months	£50	2 days	£410	2 months	£410	2 days
142	£50	6 months	£10	6 months	£410	6 months	£10	6 months
143	£100	1 week	£50	1 week	£460	1 week	£410	1 week
144	£10	6 months	£10	1 week	£10	6 months	£10	1 week
145	£200	1 week	£50	1 year	£490	1 week	£410	1 year
146	£200	1 month	£10	2 months	£490	1 month	£10	2 months
147	£500	6 months	£500	2 months	£500	6 months	£500	2 months
148	£500	6 months	£200	1 year	£500	6 months	£490	1 year
149	£500	2 days	£200	1 year	£500	2 days	£490	1 year
150	£20	1 day	£10	1 week	£310	1 day	£10	1 week

Table A.4: Experiment 4 Choices

	Positive Skew				Negative Skew			
ID	x_1	t_1	x_2	t_2	x_1	t_1	x_2	t_2
1	£200	1 year	£100	6 months	£400	10 months	£200	2 months
2	£500	6 months	£300	1 week	£200	1 year	£100	8 months
3	£400	6 months	£300	1 month	£300	10 months	£200	8 months
4	£300	2 weeks	£200	1 week	£300	2 months	£200	1 day
5	£300	1 year	£200	1 day	£500	4 months	£300	2 months

continued on next page

Table A.4 (continued from previous page)

	Positive Skew				Uniform			
ID	x_1	t_1	x_2	t_2	x_1	t_1	x_2	t_2
6	£500	2 months	£100	2 days	£300	4 months	£100	2 months
7	£500	1 month	£400	2 weeks	£200	2 months	£100	1 day
8	£300	6 months	£200	1 day	£500	6 months	£100	4 months
9	£300	2 weeks	£100	1 week	£400	1 year	£300	1 day
10	£400	2 months	£100	1 month	£300	6 months	£100	2 months
11	£500	2 months	£200	2 weeks	£300	4 months	£100	1 day
12	£500	1 year	£400	1 month	£500	8 months	£200	6 months
13	£500	2 weeks	£300	2 days	£500	8 months	£200	2 months
14	£400	1 year	£100	2 days	£500	1 year	£300	1 day
15	£500	2 months	£100	1 week	£400	6 months	£300	2 months
16	£500	1 year	£200	2 days	£500	1 year	£100	10 months
17	£400	2 weeks	£100	1 day	£500	1 year	£300	6 months
18	£300	6 months	£100	2 months	£500	10 months	£100	6 months
19	£500	2 weeks	£400	1 day	£500	8 months	£300	4 months
20	£500	6 months	£400	2 days	£300	10 months	£100	6 months
21	£400	2 months	£300	2 weeks	£300	8 months	£200	2 months
22	£500	2 weeks	£300	1 week	£400	1 year	£300	6 months
23	£500	6 months	£400	2 months	£500	4 months	£100	1 day
24	£500	1 week	£100	1 day	£300	1 year	£100	10 months
25	£400	2 months	£100	1 day	£400	10 months	£200	4 months
26	£400	1 month	£300	2 weeks	£400	1 year	£300	8 months
27	£500	2 months	£300	1 week	£500	10 months	£300	6 months
28	£400	1 month	£200	1 week	£400	8 months	£100	6 months
29	£200	6 months	£100	1 week	£500	4 months	£100	2 months
30	£200	2 weeks	£100	1 week	£300	1 year	£200	8 months
31	£500	1 week	£300	1 day	£400	8 months	£300	4 months
32	£500	6 months	£100	1 month	£400	10 months	£300	4 months
33	£300	2 weeks	£200	2 days	£500	8 months	£300	6 months
34	£400	6 months	£100	1 month	£200	6 months	£100	4 months
35	£500	1 year	£200	1 day	£300	4 months	£200	1 day
36	£400	1 year	£300	2 days	£400	1 year	£100	1 day
37	£300	1 year	£100	1 week	£200	8 months	£100	1 day
38	£300	6 months	£100	2 weeks	£400	6 months	£300	1 day
39	£400	2 weeks	£300	1 week	£500	1 year	£200	8 months
40	£500	2 months	£300	1 day	£400	2 months	£200	1 day
41	£500	2 months	£100	1 month	£300	10 months	£100	1 day
42	£500	1 year	£200	2 months	£500	1 year	£300	8 months
43	£400	1 month	£200	1 day	£400	8 months	£200	6 months
44	£500	1 year	£200	6 months	£500	10 months	£400	6 months

continued on next page

Table A.4 (continued from previous page)

	Positive Skew				Uniform			
ID	x_1	t_1	x_2	t_2	x_1	t_1	x_2	t_2
45	£400	1 year	£300	2 weeks	£200	10 months	£100	1 day
46	£400	2 months	£200	1 day	£400	10 months	£200	6 months
47	£300	2 months	£100	2 weeks	£500	4 months	£300	1 day
48	£400	2 months	£300	2 days	£500	8 months	£400	6 months
49	£500	2 days	£400	1 day	£400	6 months	£100	4 months
50	£300	1 month	£200	2 days	£400	4 months	£200	1 day
51	£500	1 month	£400	2 days	£400	1 year	£100	2 months
52	£400	1 year	£200	1 day	£500	1 year	£100	4 months
53	£500	2 weeks	£200	2 days	£200	10 months	£100	8 months
54	£400	1 month	£100	1 week	£200	8 months	£100	4 months
55	£200	2 weeks	£100	2 days	£500	8 months	£100	2 months
56	£500	1 year	£300	1 week	£500	10 months	£100	2 months
57	£400	2 days	£300	1 day	£400	8 months	£200	1 day
58	£400	1 month	£100	1 day	£400	1 year	£100	6 months
59	£300	1 month	£200	1 week	£400	8 months	£100	4 months
60	£500	1 month	£300	1 week	£500	10 months	£400	4 months
61	£200	2 months	£100	2 weeks	£400	10 months	£300	8 months
62	£400	2 months	£300	1 day	£400	1 year	£200	8 months
63	£400	1 month	£300	2 days	£200	4 months	£100	1 day
64	£400	6 months	£300	1 week	£400	10 months	£300	2 months
65	£200	2 weeks	£100	1 day	£500	10 months	£300	1 day
66	£300	1 month	£100	2 weeks	£500	6 months	£200	4 months
67	£300	1 week	£100	1 day	£400	2 months	£300	1 day
68	£500	2 months	£200	1 month	£400	10 months	£100	2 months
69	£500	2 weeks	£200	1 day	£500	6 months	£100	2 months
70	£400	1 year	£100	2 weeks	£400	8 months	£300	2 months
71	£200	1 week	£100	2 days	£300	6 months	£100	4 months
72	£400	2 weeks	£200	2 days	£200	4 months	£100	2 months
73	£500	6 months	£300	1 day	£500	6 months	£300	1 day
74	£200	2 months	£100	1 week	£500	8 months	£300	2 months
75	£400	1 week	£200	1 day	£400	6 months	£100	2 months
76	£300	1 year	£200	2 weeks	£400	4 months	£300	1 day
77	£500	1 month	£200	2 days	£500	10 months	£200	6 months
78	£200	1 year	£100	1 week	£500	8 months	£200	1 day
79	£300	2 weeks	£100	1 day	£400	8 months	£200	4 months
80	£400	1 year	£200	6 months	£500	2 months	£300	1 day
81	£400	2 months	£100	2 weeks	£500	8 months	£200	4 months
82	£400	1 year	£300	1 week	£200	8 months	£100	6 months
83	£400	2 months	£300	1 week	£500	1 year	£200	2 months

continued on next page

Table A.4 (continued from previous page)

	Positive Skew				Uniform			
ID	x_1	t_1	x_2	t_2	x_1	t_1	x_2	t_2
84	£500	1 year	£100	1 month	£500	10 months	£200	1 day
85	£200	1 month	£100	2 weeks	£300	6 months	£200	2 months
86	£200	2 months	£100	2 days	£500	10 months	£400	2 months
87	£200	6 months	£100	1 month	£500	6 months	£200	2 months
88	£500	1 year	£400	1 week	£500	1 year	£100	6 months
89	£400	1 year	£200	2 months	£500	10 months	£200	8 months
90	£500	2 weeks	£200	1 week	£500	6 months	£400	1 day
91	£500	1 year	£400	1 day	£500	8 months	£400	2 months
92	£400	2 weeks	£200	1 week	£300	1 year	£200	6 months
93	£500	2 months	£200	1 day	£400	8 months	£300	1 day
94	£400	2 weeks	£300	2 days	£500	10 months	£100	4 months
95	£400	6 months	£300	2 months	£400	8 months	£100	1 day
96	£300	1 year	£200	1 week	£400	1 year	£300	4 months
97	£300	6 months	£200	2 days	£200	10 months	£100	6 months
98	£400	6 months	£300	1 day	£300	8 months	£100	6 months
99	£500	1 month	£100	1 week	£500	1 year	£100	1 day
100	£300	1 year	£200	2 months	£300	1 year	£100	6 months
101	£400	1 year	£100	1 month	£500	1 year	£400	8 months
102	£300	1 month	£100	1 day	£400	6 months	£300	4 months
103	£500	1 month	£300	1 day	£400	8 months	£300	6 months
104	£400	6 months	£200	1 month	£500	8 months	£400	4 months
105	£200	6 months	£100	2 months	£300	10 months	£200	1 day
106	£400	1 week	£300	1 day	£400	10 months	£100	1 day
107	£300	1 month	£200	1 day	£500	10 months	£400	8 months
108	£400	1 year	£100	1 week	£300	1 year	£200	2 months
109	£500	1 year	£200	1 week	£500	1 year	£200	4 months
110	£400	6 months	£300	2 days	£400	1 year	£200	4 months
111	£500	1 year	£300	6 months	£500	2 months	£400	1 day
112	£400	6 months	£100	2 months	£500	6 months	£300	4 months
113	£500	6 months	£200	2 days	£300	4 months	£200	2 months
114	£400	2 months	£100	2 days	£200	6 months	£100	2 months
115	£500	2 days	£100	1 day	£500	1 year	£300	10 months
116	£200	1 year	£100	2 months	£300	8 months	£100	4 months
117	£500	1 month	£100	1 day	£300	8 months	£100	2 months
118	£500	6 months	£200	2 weeks	£400	10 months	£200	1 day
119	£200	2 days	£100	1 day	£500	1 year	£200	10 months
120	£300	2 months	£100	1 week	£400	6 months	£200	1 day
121	£400	2 weeks	£100	1 month	£400	8 months	£300	10 months
122	£200	2 weeks	£100	6 months	£400	8 months	£200	8 months

continued on next page

Table A.4 (continued from previous page)

	Positive Skew				Uniform			
ID	x_1	t_1	x_2	t_2	x_1	t_1	x_2	t_2
123	£500	1 day	£300	2 months	£400	2 months	£200	6 months
124	£300	1 day	£200	6 months	£500	1 day	£400	6 months
125	£100	2 weeks	£100	6 months	£400	1 day	£100	1 day
126	£400	2 days	£400	6 months	£100	1 day	£100	8 months
127	£500	1 day	£300	6 months	£500	6 months	£300	10 months
128	£500	2 weeks	£200	1 year	£400	1 day	£100	1 year
129	£200	1 week	£100	1 week	£500	8 months	£300	10 months
130	£300	1 week	£100	6 months	£500	1 day	£400	1 day
131	£300	1 day	£300	2 days	£500	2 months	£300	4 months
132	£500	2 days	£400	6 months	£200	4 months	£100	10 months
133	£300	1 day	£100	2 weeks	£500	1 day	£200	4 months
134	£400	2 days	£300	1 month	£100	2 months	£100	4 months
135	£400	2 weeks	£300	1 month	£500	4 months	£100	6 months
136	£400	1 day	£200	2 days	£400	2 months	£400	4 months
137	£200	1 day	£200	2 weeks	£100	1 day	£100	1 year
138	£400	1 week	£300	2 weeks	£300	6 months	£200	8 months
139	£500	2 days	£500	6 months	£300	2 months	£200	2 months
140	£200	2 months	£200	1 year	£500	10 months	£300	1 year
141	£500	2 days	£100	1 week	£500	10 months	£100	1 year
142	£200	1 day	£100	1 day	£300	1 day	£200	6 months
143	£500	2 weeks	£400	2 months	£300	4 months	£200	8 months
144	£500	1 year	£200	1 year	£500	8 months	£200	1 year
145	£200	2 weeks	£100	2 weeks	£300	1 day	£300	4 months
146	£200	1 day	£100	2 weeks	£200	2 months	£200	10 months
147	£500	2 days	£500	1 year	£500	4 months	£300	1 year
148	£300	2 weeks	£100	6 months	£500	1 day	£500	2 months
149	£300	1 week	£200	6 months	£400	10 months	£300	10 months
150	£500	2 months	£200	1 year	£500	4 months	£400	4 months

REFERENCES

ABDELLAOUI, M. (2000): "Parameter free elicitation of utilities and probability weighting functions," *Management Science*, 46, 1497–1512.

AINSLIE, G. (1975): "Specious reward: A behavioural theory of impulsiveness and impulse control," *Psychological Bulletin*, 82, 463–496.

ALLAIS, M. (1953): "Le comportement de l'homme rationel devant le risque: Critique des postulats et axioms de l'ecole americaine [Rational man's behavior in face of risk: Critique of the American School's postulates and axioms]," *Econometrica*, 21, 503–546.

ARIELY, D., B. KŐSZEGI, N. MAZAR, AND K. SHAMPAN'ER (2008): "Price-sensitive preferences," .

BECKER, G. M., M. H. DEGROOT, AND J. MARSCHAK (1963): "Stochastic models of choice behavior," *Behavioral Science*, 8, 41–55.

——— (1964): "Measuring utility by a single-response sequential method," *Behavioral Science*, 9, 226–232.

BELL, D. E. (1985): "Disappointment in decision making under uncertainty," *Operations Research*, 33, 1–27.

BENARTZI, S., AND R. H. THALER (2001): "Naive diversification strategies in defined contribution saving plans," *American Economic Review*, 91, 79–98.

BENTHAM, J. (1970): *An introduction to the principles of morals and legislation*. Athlone, London.

BIRNBAUM, M. H. (1992): "Violations of the monotonicity and contextual effects in choice-based certainty equivalents," *Psychological Science*, 3, 310–314.

——— (2008): "New paradoxes of risky decision making," *Psychological Review*, 115, 453–501.

BIRNBAUM, M. H., AND A. CHAVEZ (1997): "Tests of theories of decision making: Violations of branch independence and distribution independence," *Organizational Behavior and Human Decision Processes*, 71, 161–194.

BLAVATSKYY, P. R. (2007): "Stochastic expected utility theory," *Journal of Risk and Uncertainty*, 34, 259–286.

BLAVATSKYY, P. R., AND G. POGREBNA (2010): "Models of stochastic choice and decision theories: Why both are important for analyzing decisions," *Journal of Applied Econometrics*, 25, 963–986.

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.

BOYCE, C. J. (2009): "Subjective well-being: An intersection between economics and psychology," Ph.D. thesis, Warwick University, Coventry, England.

BRANDSTÄTTER, E., G. GIGERENZER, AND R. HERTWIG (2006): "The priority heuristic: Making choices without trade-offs," *Psychological Review*, 113, 409–432.

BROWN, G. D. A., J. GARDNER, A. J. OSWALD, AND J. QIAN (2008): "Does wage rank affect employees' well-being?," *Industrial Relations*, 47, 355–389.

BUSEMEYER, J. R. (1985): "Decision making under uncertainty: A comparison of simple scalability, fixed sample, and sequential sampling models," *Journal of Experimental Psychology*, 11, 538564.

BUSEMEYER, J. R., AND J. T. TOWNSEND (1993): "Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment," *Psychological Review*, 100, 432–459.

CAMERER, C. F. (1995): "Individual decision making," in *Handbook of experimental economics*, ed. by J. Kagel, and A. E. Roth, pp. 587–703. Princeton University Press, Princeton, NJ.

DORLET, A., I. SIMONSON, AND A. TVERSKY (2000): "Indifference curves that travel with the choice set," *Marketing Letters*, 11, 199–209.

EDWARDS, W. (1962): "Subjective probabilities inferred from decisions," *Psychological Review*, 69, 109–135.

EREV, I., A. E. ROTH, R. L. SLONIM, AND G. BARRON (2002): "Combining a theoretical prediction with experimental evidence," .

GONZALEZ, R., AND G. WU (1999): "On the shape of the probability weighting function," *Cognitive Psychology*, 38, 129–166.

GUL, F. (1991): "A theory of disappointment aversion," *Econometrica*, 59, 667–686.

KAHNEMAN, D., J. L. KNETSCH, AND R. H. THALER (1991): "Anomalies: The endowment effect, loss aversion, and status quo bias," *Journal of Economic Perspectives*, 5, 193–206.

Kahneman, D., and A. Tversky (1979): "Prospect theory: An analysis of decision under risk," *Econometrica*, 47, 263–291.

Kornienko, T. (2011): "A cognitive basis for context-dependent utility," .

Kőszegi, B., and M. Rabin (2006): "A model of reference-dependent preferences," *Quarterly Journal of Economics*, 121, 1133–1165.

——— (2007): "Reference-dependent risk attitudes," *American Economic Review*, 97, 1047–1073.

Laibson, D. (1997): "Golden eggs and hyperbolic discounting," *Quarterly Journal of Economics*, 112, 443–477.

Loewenstein, G., and D. Prelec (1992): "Anomalies in intertemporal choice: Evidence and an interpretation," *Quarterly Journal of Economics*, 107, 573–597.

Loomes, G. (2005): "Modelling the stochastic component of behaviour in experiments: Some issues for the interpretation of data," *Experimental Economics*, 8, 301–323.

——— (2010): "Modeling choice and valuation in decision experiments," *Psychological Review*, 117, 902–924.

Loomes, G., and R. Sugden (1982): "Regret theory: An alternative theory of rational choice under uncertainty," *Economic Journal*, 92, 805–824.

——— (1986): "Disappointment and dynamic consistency in choice under uncertainty," *Review of Economic Studies*, 53, 271–282.

Luce, R. D. (1959): *Individual choice behavior*. Wiley, New York.

Luce, R. D. (2000): *Utility of gains and losses: Measurement-theoretical and experimental approaches*. Erlbaum, Mahwah, NJ.

Maccheroni, F., M. Marinacci, and A. Rustichini (2009a): "Pride and diversity in social economics," .

——— (2009b): "Social decision theory: Choosing within and between groups," .

Mazur, J. E. (1987): "An adjusting procedure for studying delayed reinforcement," in *Quantitative analyses of behavior: Vol. 5. The effect of delay and of intervening events on reinforcement value*, ed. by M. L. Commons, J. E. Mazur, J. A. Nevin, and H. Rachlin, pp. 55–73. Erlbaum, Hillsdale, NJ.

Parducci, A. (1965): "Category judgment: A range-frequency model," *Psychological Review*, 72, 407–418.

——— (1995): *Happiness, pleasure and judgment: The contextual theory and its applications*. Erlbaum, Mahwah, NJ.

Prelec, D. (1998): "The probability weighting function," *Econometrica*, 66, 497–527.

Quiggin, J. (1993): *Generalized expected utility theory: The rank-dependent model*. Kluwer Academic Publishers, Norwell, MA.

Rachlin, H., A. Raineri, and D. Cross (1991): "Subjective probability and delay," *Journal of the Experimental Analysis of Behavior*, 55, 233–244.

Read, D. (2001): "Is time-discounting hypoerbolic or subadditive?," *Journal of Risk and Uncertainty*, 23, 5–32.

——— (2007): "Experienced utility: Utility theory from Jeremy Bentham to Daniel Kahneman," *Thinking & Reasoning*, 13, 45–61.

Samuelson, P. A. (1937): "A note on measurement of utility," *Review of Economic Studies*, 4, 155–161.

Savage, L. J. (1954): *The foundations of statistics*. Wiley, New York.

Schoemaker, P. J. H. (1982): "The expected utility model: Its variants, purposes, evidence, and limitations," *Journal of Economic Literature*, 20, 529–563.

Scholten, M., and D. Read (2010): "The psychology of intertemporal tradeoffs," *Psycho-*

logical Review, 117, 925–944.

SEYMOUR, B., AND S. M. MCCLURE (2008): “Anchors, scales and the relative coding of value in the brain,” *Current Opinion in Neurobiology*, 18, 173–178.

SHEPARD, R. N. (1957): “Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space,” *Psychometrica*, 22, 325–345.

STARMER, C. (2000): “Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk,” *Journal of Economic Literature*, 38, 332–382.

STEWART, N. (2009): “Decision by sampling: The role of the decision environment in risky choice,” *Quarterly Journal of Experimental Psychology*, 62, 1041–1062.

STEWART, N., N. CHATER, AND G. D. A. BROWN (2006): “Decision by sampling,” *Cognitive Psychology*, 53, 1–26.

STEWART, N., N. CHATER, H. P. STOTT, AND S. REIMERS (2003): “Prospect relativity: How choice options influence decision under risk,” *Journal of Experimental Psychology: General*, 132, 23–46.

STEWART, N., AND K. SIMPSON (2008): “A decision-by-sampling account of decision under risk,” in *The probabilistic mind: Prospects for Bayesian cognitive science*, ed. by N. Chater, and M. Oaksford, pp. 261–276. Oxford University Press, Oxford, England.

STIGLER, G. J. (1950a): “The development of utility theory. I,” *Journal of Political Economy*, 58, 307–327.

——— (1950b): “The development of utility theory. II,” *Journal of Political Economy*, 58, 373–396.

THALER, R. H. (1981): “Some empirical evidence on dynamic inconsistency,” *Economic Letters*, 8, 201–207.

TOM, S. M., C. R. FOX, C. TREPEL, AND R. A. POLDRACK (2007): “The neural basis of loss aversion in decision-making under risk,” .

TREMBLAY, L., AND W. SCHULTZ (1999): “Relative reward preference in primate orbitofrontal cortex,” .

TVERSKY, A., AND D. KAHNEMAN (1992): “Advances in prospect theory: Cumulative representation of uncertainty,” *Journal of Risk and Uncertainty*, 5, 297–323.

UNGEMACH, C., N. STEWART, AND S. REIMERS (2011): “How incidental values from our environment affect decisions about money, risk, and delay,” *Psychological Science*, 22, 253–260.

VON NEUMANN, M., AND O. MORGENSTERN (1947): *Theory of games and economic behavior*. Princeton University Press, Princeton, NJ, 2nd edn.

WU, G., AND R. GONZALEZ (1996): “Curvature of the probability weighting function,” *Management Science*, 42, 1676–1690.

——— (1999): “Nonlinear decision weights in choice under uncertainty,” *Management Science*, 45, 74–85.