



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

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

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

To cite this article:

Neil Stewart, Stian Reimers, Adam J. L. Harris (2015) On the Origin of Utility, Weighting, and Discounting Functions: How They Get Their Shapes and How to Change Their Shapes. *Management Science* 61(3):687-705. <http://dx.doi.org/10.1287/mnsc.2013.1853>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2015, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

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

Neil Stewart

Department of Psychology, University of Warwick, Coventry CV4 7AL, United Kingdom,
neil.stewart@warwick.ac.uk

Stian Reimers

Department of Psychology, City University London, London EC1V 0HB, United Kingdom,
stian.reimers@city.ac.uk

Adam J. L. Harris

Department of Cognitive, Perceptual and Brain Sciences, University College London,
London WC1H 0AP, United Kingdom, adam.harris@ucl.ac.uk

We present a theoretical account of the origin 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 choices used to elicit them. The data demonstrate that there is no stable mapping between attribute values and their subjective equivalents. Expected and discounted utility theories, and also their descendants such as prospect theory and hyperbolic discounting theory, simply assert stable mappings to describe choice data and offer no account of the instability we find. We explain where the shape of the mapping 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.

Data, as supplemental material, are available at <http://dx.doi.org/10.1287/mnsc.2013.1853>.

Keywords: 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

History: Received September 5, 2012; accepted September 19, 2013, by Yuval Rottenstreich, judgment and decision making. Published online in *Articles in Advance* March 18, 2014.

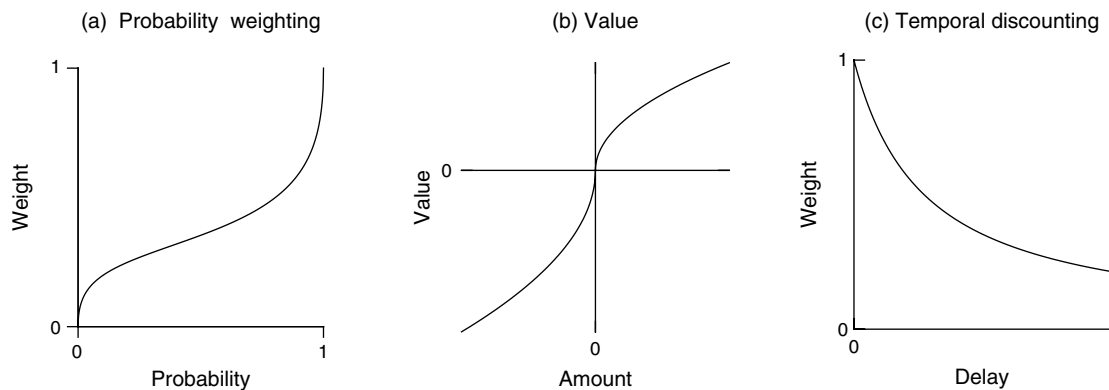
1. Introduction

Central to our economic behavior are the attributes money, probability, and time. Our representations of these attributes are thought to determine our economic behavior. 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 choice that 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. This means that shape of utility, weighting, and discounting functions is a property of the choice environment and not just of the individual. We argue that although it is possible to derive utility, weighting, and discounting functions from behavioral data (such as a series of choices), these functions do not describe the process of choosing. Our theoretical account of the choice process does not involve utility, weighting, and discounting functions, but does explain both the success of the traditional approach and also accounts for the systematic variation of these functions across question sets.

2. Decision Under Risk

Expected utility is the normative model for risky decision making (von Neumann and Morgenstern 1947). In expected utility theory, wealth is translated into utility by a utility function. To choose between risky

Figure 1 Psychoeconomic Functions Convert (a) Probability, (b) Amount, and (c) Delay Into Their Subjective Equivalents

options, the average or expected utility of each available course of action is calculated and the action maximizing expected utility is selected. Expected utility theory is often paired with the assumption of 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, a concave utility function is used, which is 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.

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, Bussemeyer 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 decision weight. Figures 1(a) and 1(b) give example weighting and value functions.

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 (for reviews, see Allais 1953, Birnbaum 2008, Camerer 1995, Loomes 2010, Luce 2000, Schoemaker 1982, Starmer 2000), but here we use just two—the finding that, in choices about gains, people are risk averse for gambles involving medium and large probabilities, but risk seeking for those involving small probabilities—to illustrate how the shapes of the value and weight functions were constructed to account for choice data. Although Tversky and Kahneman (1992) 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, designed to fit particular choice patterns. Consider the shape of the value function for gains (the top right quadrant of Figure 1(b)). In an initial choice between (A) 3,000 for sure and (B) an 80% chance of 4,000 otherwise nothing, people are risk averse and prefer A. In prospect theory, because the value function is concave, the value of 3,000 is a bit more than 80% of the value of 4,000, and so people prefer the sure 3,000; that is, the value function is concave for gains to describe risk aversion.

The shape of the weighting function is also constructed to describe the choices that people make. In a choice between (C) a 25% chance of 3,000 otherwise nothing and (D) a 20% chance of 4,000 otherwise nothing, people are risk seeking and prefer D. Because the C–D choice is derived from the A–B choice by multiplying probabilities by 1/4, participants should prefer the A and C or prefer B and D. Their preference for A and D, known as the common ratio effect (Allais 1953), is not consistent with expected utility theory. Kahneman and Tversky (1979) account for these data using the weighting function, in which 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 80% and 100% are quite different, but the weights of 20% and 25% are quite similar (see Figure 1(a)); that is, the weighting function takes its inverse-S shape 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). Here we select one example behavioral finding to illustrate the descriptive model, but there are many (for reviews, see Loewenstein and Prelec 1992, Scholten and Read 2010). 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 forward 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(c)). 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.

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 et al. (2011) examined how the distribution of attribute values we experience every day affects choices. In one study, Ungemach et al. (2011) 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 et al. (2011) 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 (see Matthews 2012 for a failure to replicate the birthday experiment and Peetz and Wilson 2013 for a successful study using birthdays as temporal landmarks). 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. Beauchamp et al. (2012), Birnbaum (1992), and Stewart et al. (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 (1992) found that the

valuation of a target gamble was higher when people selected a valuation from a negatively skewed set than from a positively skewed set. If perceived value is, in part, determined by rank position, a particular valuation will seem larger in the positively skewed set (where it has a high rank position) than in a negatively skewed set (where it has a low rank position). Thus, the value of the target gamble is matched to a lower valuation in the positively skewed set (because the valuations seem higher) and a higher valuation in the negatively skewed set (because the valuations seem lower). Using an incentive compatible auction, Mazar et al. (2014) 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). An example of range affecting valuations is given by Vlaev et al. (2009). They used a Becker–DeGroot–Marschak auction to show that doubling the range of prices from which a sale price was to be drawn doubled the price people would pay to avoid a series of electric shocks. Haggag and Paci (2013) report a similar effect in a natural experiment where taxi customers give higher tips when offered 20%, 25%, and 30% as potential defaults compared to 15%, 20%, and 25%—there could be other causes for this taxi tip effect, but it is consistent with the rank hypothesis.

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. Benartzi and Thaler (2001) account for this pattern with the $1/n$ heuristic, which assumes people place equal funds into each option, but the allocation is also compatible with the rank hypothesis. Stewart et al. (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

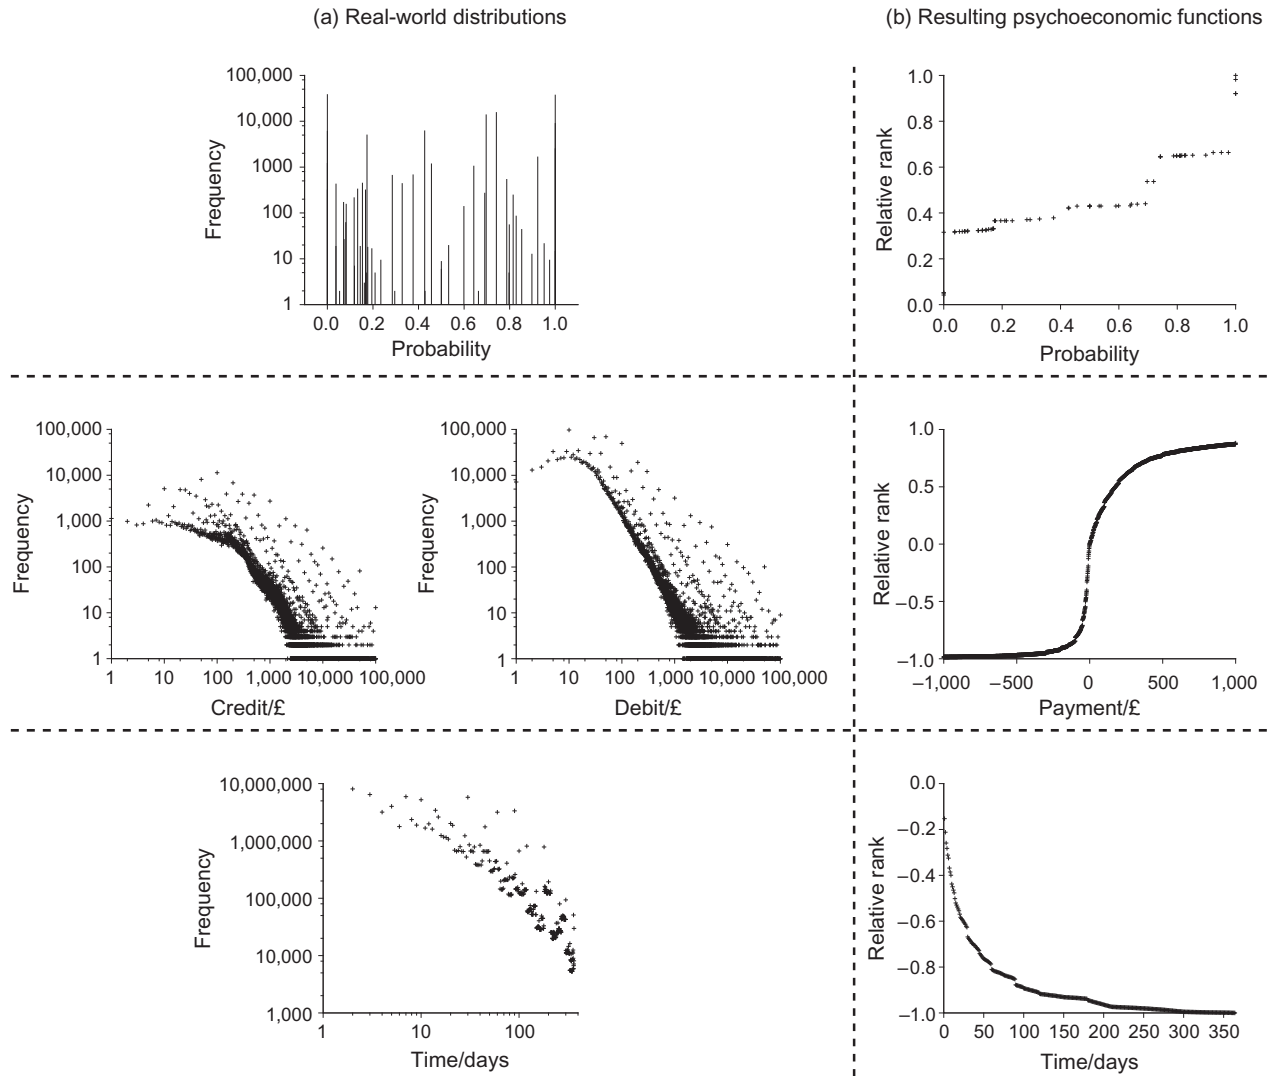
delay) is determined, at least in part, by its rank position in the set of values currently in a person's head. So, for example, \$10 has a higher subjective value in the set \$2, \$5, \$8, and \$15 because it ranks 2nd, but has a lower subjective value in the set \$2, \$15, \$19, and \$25 because it 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 (Brown et al. 2008, Boyce 2009).

Stewart et al. (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, and \$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 et al. (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. Assuming that people's memories reflect the world in which they live, Stewart et al. (2006) used these real-world distributions as approximations of the contents of people's memories and estimated the value, weighting, and discounting functions that would be observed under decision by sampling (right column). The top left panel is the distribution of probability phrases occurring in everyday English. People prefer to use verbal phrases rather than numbers. In a new analysis, we find that of the first 500 sentences in the British National Corpus containing the six-letter string "chance," 65 are about the probability of an event and, of these, 60 use verbal phrases and 5 use numbers. And for responses to the open question of Ungemach et al. (2011, p. 256), "How likely do you think it is to rain tomorrow?" 340 involved verbal phrases and 47 involved numbers. The top right panel is the resulting weighting function expected if people compare

Figure 2 Real-World Attribute-Value Distributions (a) and the Resulting Psychoeconomic Functions (b) for Probability (Top), Money (Middle), and Delay (Bottom)



Source. Adapted from Stewart et al. (2006, Figures 1, 2, 5, 6, 7).
 Note. See the text for details.

probabilities against those they encounter in everyday English. The function is inverse S-shaped because high- and low-probability phrases are more frequent than intermediate-probability phrases. The function is asymmetrical because high-probability phrases are more frequent than low-probability phrases. The middle row repeats the exercise for value. The real-world distributions on the left are for credits and debits to bank accounts and are positively skewed. The right panel shows the resulting value functions if people compare gains and losses against these credits and debits. Notice how, because there are more small gains in the world than large gains, the value function is concave for gains: Against the distribution of gains, a £100 increase in a prize from £100 to £200 improves the rank position substantially more than a £100 increase from £900 to £1,000. The bottom row

repeats the exercise for time. The left panel shows the frequency with which different delays occur in text on the Internet. The right panel shows the resulting weighting function if people compare delays to the real-world distribution on the Internet. (Because larger losses and delays are worse than shorter losses and delays, they have been plotted with negative ranks.) 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 the decision-by-sampling 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 online decision-by-sampling calculator, which gives exact choice probabilities for any

set of prospects, is available at <http://www.stewart.warwick.ac.uk>.

Equation (1) formalizes the expression for the subjective value of an attribute, although this is hardly 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 $Y = \{y_1, y_2, \dots, y_n\}$ is given by

$$s(x, Y) = \frac{\sum_{y \in Y} c(x, y)}{n}, \quad (1)$$

where

$$c(x, y) = \begin{cases} 1 & \text{if } x \text{ compares favorably to } y, \\ 0 & \text{if } x \text{ does not compare favorably to } y. \end{cases} \quad (2)$$

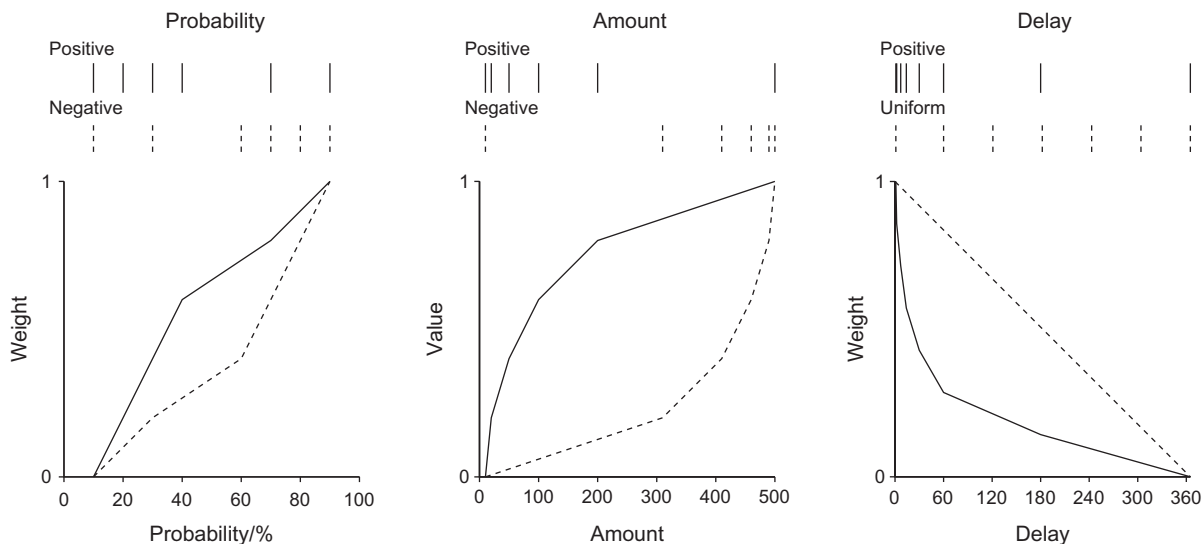
The subjective value $s(x, Y)$ is the probability of a favorable comparison between x and a randomly selected member of set Y . Equivalently, $s(x, Y)$ is the rank of x in Y , normalized to lie between zero and one.

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 some of 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 (center 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, as 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.

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 et al. (2006), we know that the distributions of money in the world are positively skewed, with small amounts being highly frequent

Figure 3 Hypothetical Psychoeconomic Functions for Probability (Left), Amount (Center), and Delay (Right)



Note. The subjective value of an attribute is its rank within the distribution of other attribute values (shown above each plot), normalized to lie between zero and one.

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 systematically and predictably 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 not to a person's stable risk or intertemporal preferences, but instead, at least in part, to the experimenter's choice of attribute values used in the experiment.

6. Experiments

6.1. Experiment 1A

Participants made choices of the form “ p chance of x , otherwise nothing” or “ q chance of y , otherwise nothing.” Each choice was between a smaller probability of a larger amount of money or a larger probability of a smaller amount of money. Between participants, the distribution of amounts available was manipulated to be either positively skewed or negatively skewed.

6.1.1. Method. Participants. Forty-one Warwick psychology first-year undergraduates participated. Course credit was given for attending. In addition, for each participant, the gambles they selected on two randomly sampled choices were played via urn draws. After applying an experiment exchange rate to any winnings, 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 the others in this paper, made no catch-trial errors.

Design. A set of five probabilities was crossed with a set of six amounts to create 30 gambles of the form “ p chance of x , otherwise nothing.” All participants experienced probabilities 0.2, 0.4, 0.6, 0.8, and 1.0. Participants were randomly assigned to receive either a positively or a negatively skewed set of amounts (see Figure 3, center 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 used to create a set of all possible pairwise choices between them. 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. Choices and instructions for all experiments are available at <http://www.stewart.warwick.ac.uk>.

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 nonwinning 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 for their selected gambles to determine their winnings. 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.

Note that in all of the experiments reported here, payments were incentive compatible (except Experiment 1C), and there was no deception.

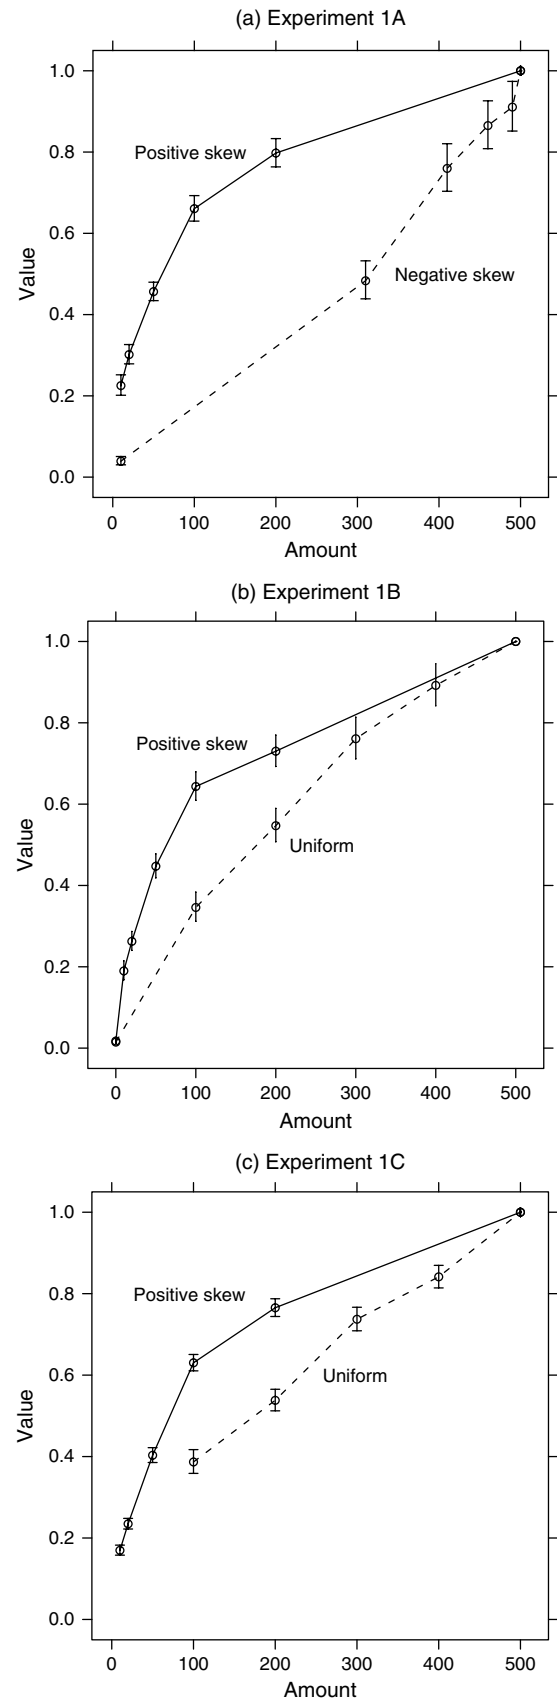
6.1.2. Results and Discussion. For each participant, raw data were the 150 choices between a relatively safe gamble offering a “ q chance of y , otherwise nothing” and a relatively risky gamble offering a “ p chance of x , otherwise nothing,” where $q > p$ and $y < x$. To recover the subjective values of the amounts of money, we fitted Equation (3) to the choice data:

$$\text{Prob}(\text{safe}) = \frac{\text{bias}_{\text{cond}} [q u_{\text{cond}}(y)]^{\gamma_{\text{cond}}}}{\text{bias}_{\text{cond}} [q u_{\text{cond}}(y)]^{\gamma_{\text{cond}}} + [p u_{\text{cond}}(x)]^{\gamma_{\text{cond}}}}, \quad (3)$$

where $q u_{\text{cond}}(y)$ is the expected utility of the safe gamble in condition cond , and $p u_{\text{cond}}(x)$ is the expected utility of the risky gamble in condition cond . This Luce (1959)–Shepard (1957) choice formulation is a straightforward way to incorporate a stochastic component in the expected utility model, giving an increasing probability of selecting the safe gamble as the utility of the safe gamble increases or the utility of the risky gamble decreases. The γ parameter controls the degree of determinism in the model: $\gamma = 1$ gives choice probabilities proportional to the expected utilities, and $\gamma > 1$ gives more extreme choice probabilities, so gambles with only slightly higher expected utility are very likely to be chosen. The bias parameter allows for an overall bias toward safe or risky choices independently of the amounts and probabilities in the gambles. (Though we do find bias differences in some experiments, fixing bias across conditions produces almost identical revealed utility, weighting, and discounting functions and the same pattern of statistical significance in for every experiment in this paper, except in two places we note below.) The cond subscripts indicate that the utility function $u(\cdot)$ can differ between conditions, as can the γ and bias parameters. Appendix A explains how $u_{\text{cond}}(\text{amount}_i)$ was estimated for each amount_i in each condition as part of an off-the-shelf mixed-effects logistic regression and reports the full set of estimated parameters.

The revealed utility functions that maximize the likelihood of the choice data are shown in Figure 4(a).

Figure 4 Utility Functions for Experiments 1A–1C



Note. Error bars are standard errors.

Without loss of generality, the utility of £500 was set at 1. Effectively, the logistic regression adjusted the heights of each point in the utility functions to best fit the choices participants made. Though they were not constrained to do so, functions increase monotonically. 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. As a summary test of the difference in concavity, we found that the best fitting power α_{cond} for utility functions $u_{cond}(x) = x^{\alpha_{cond}}$ differ significantly across conditions ($\chi^2(1) = 6.36, p = 0.012$). Figure B.1 in Appendix B reports estimates for *bias*, γ , and α for each condition and the differences between conditions with bootstrapped 95% confidence intervals and shows significant α differences for this experiment and the later experiments. And this is the core result: Participants who experienced a different distribution of amounts chose as if they had a different shaped utility function. In particular, participants in the positive-skew condition value £200 more than participants in the negative-skew condition value £310, even though £200 is smaller than £310 ($\chi^2(1) = 7.05, p = 0.0079$).

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 et al. 2011).

6.2. Experiment 1B

Experiments 1B and 1C are systematic replications of Experiment 1A. Both follow the design of Experiment 1A, except for the differences noted here.

6.2.1. Method. Participants. Fifty-six first-year undergraduates participated in Experiment 1B. Each experiment in this paper used a different set of participants. For each participant, the gamble they selected on one randomly sampled choice was played via an urn draw. After applying an experiment exchange rate to any winnings, participants could win up to 5. No participants violated stochastic dominance on more than 10% of catch trials, so all data were retained.

Design. Participants were randomly assigned to receive either a positive-skew condition with amounts £0, £10, £20, £50, £100, £200, and £500, or uniform condition with amounts £0, £100, £200, £300, £400, and £500. Otherwise gambles and choices were constructed as in Experiment 1A. The uniform condition replaces the negative-skew condition from

Experiment 1A to make the study more representative of distributions used in other experiments and so that £100 and £200 are common across conditions.

Procedure. This replication omitted the panel of attribute values listing the amounts and probabilities on offer, which one reviewer was concerned might be driving the effects. Instead participants only saw the series of choices, so any effects of distribution would have to be the effect of the attributes viewed on earlier choices affecting later choices.

6.2.2. Results and Discussion. Figure 4(b) shows a concave utility function for a positive-skew condition and a linear utility function for a uniform condition. The difference in concavity is significant ($\chi^2(1) = 6.99, p = 0.0082$). In this experiment, amounts £100 and £200 are common to both distributions, but the revealed values are higher in the positive-skew condition because their rank position in the distribution is higher ($\chi^2(1) = 26.96, p < 0.0001$ for the difference at £100, and $\chi^2(1) = 7.16, p = 0.0074$ for the difference at £200).

6.3. Experiment 1C

6.3.1. Method. Experiment 1C replicates Experiments 1A and 1B using an online sample of UK survey participants making hypothetical rather than incentivized choices. One hundred twenty-four participants completed the experiment in their Web browsers. Data from 17 participants were deleted for violating dominance on more than 10% of catch trials, though these deletions do not alter the pattern of results. Conditions were as Experiment 1B except the £0 amounts were dropped.

6.3.2. Results and Discussion. Figure 4(c) plots the utility functions recovered for the positive-skew and uniform conditions, replicating the concave and linear utility functions from Experiment 1B. Fits of power law utility functions shows the difference in concavity is marginal ($\chi^2(1) = 3.50, p = 0.06$). (This marginal result is significant when *bias* is held constant across conditions.) There are significant differences in values at £100, $\chi^2(1) = 59.79$ and $p < 0.0001$, and £200, $\chi^2(1) = 50.47$ and $p < 0.0001$.

Finally, we note that Mullett (2012) has replicated Experiment 1A using a within-participants manipulation where the prices of different categories of product (holidays versus mobile phones) differed in skew, and Matthews (2013) has replicated Experiment 1A using a within-participants manipulation where two different gamble sets with different amount distributions were offered by different experimenters.

6.4. Experiment 2A

In Experiment 2A, 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 Experiments 1A–1C.

6.4.1. Method. 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. 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. One hundred twenty nonstochastically 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 participants understood the probabilities and the method for resolving them. In this study, probabilities were resolved by drawing 1 of 100 numbered 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 × 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.

6.4.2. Results and Discussion. The analysis repeats the procedure from Experiment 1, except subjective probabilities instead of utilities varied between conditions; that is, we fitted

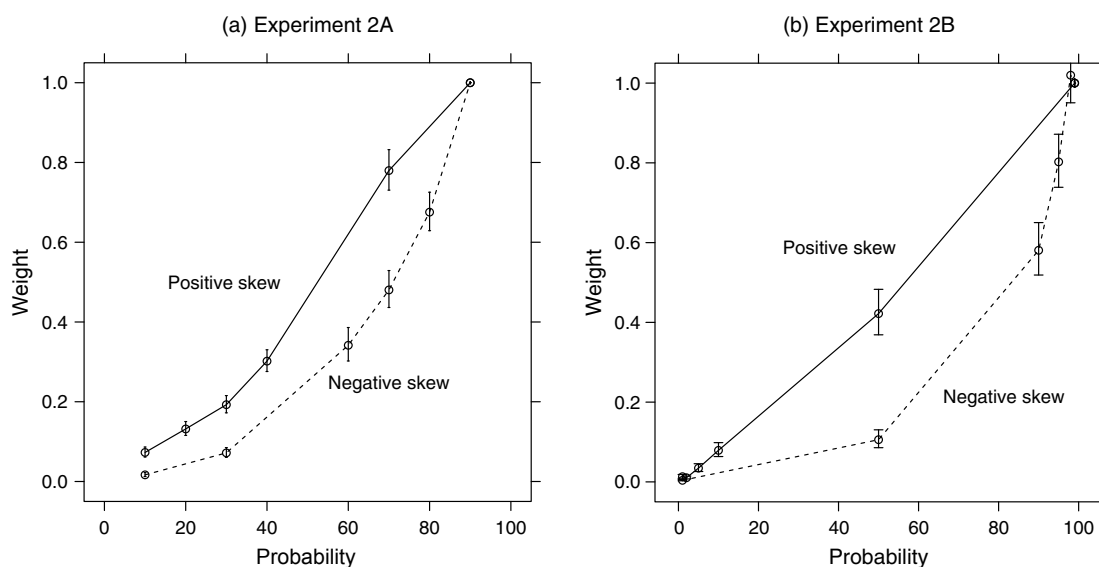
$$\text{Prob}(safe) = \frac{bias_{cond}[w_{cond}(q)y]^{\gamma_{cond}}}{bias_{cond}[w_{cond}(q)y]^{\gamma_{cond}} + [w_{cond}(p)x]^{\gamma_{cond}}}, \quad (4)$$

where $w_{cond}(p)$ gives the weighting of probability p in condition $cond$. Figure 5(a) shows the revealed weighting functions. With weighting functions constrained to follow power laws, the test of a difference in concavity did not reach significance ($\chi^2(1) = 2.18$, $p = 0.13$). However, modeling with free weightings revealed significant differences across conditions for the common 30% ($\chi^2(1) = 18.18$, $p < 0.0001$) and 70% ($\chi^2(1) = 14.31$, $p = 0.0002$).

6.5. Experiment 2B

6.5.1. Method. Experiment 2B is a systematic replication of Experiment 2A with more highly skewed distributions, and differs only as noted here. Thirty-six Warwick psychology first-year undergraduates

Figure 5 Utility Functions for Experiments 2A and 2B



Note. Error bars are standard errors.

participated. The positive-skew condition used probabilities 1%, 2%, 5%, 10%, 50%, and 99%. The negative-skew condition used probabilities 1%, 50%, 90%, 95%, 98%, and 99%. No participant violated dominance in more than 10% of catch trials, so all data were retained.

6.5.2. Results and Discussion. Figure 5(b) shows that the weighting function in the positive-skew condition is more concave than the negative-skew condition. Modeling using power law weighting functions reveals a significant difference between conditions ($\chi^2(1) = 181.5$, $p < 0.0001$). Modeling with free weightings tested whether the weighting for the common probability 50% differed across conditions. The weighting of 50% is significantly higher in the positive-skew condition ($\chi^2(1) = 41.72$, $p < 0.0001$), consistent with the rank hypothesis.

6.6. Experiment 3A

In Experiments 3A, 3B, and 4, we repeat Experiments 1 and 2, but using delay instead of risk. Participants made choices of the form “ x at time t ” or “ y at time u .” Each choice was between a later larger amount and a smaller sooner amount. In Experiments 3A and 3B, the distribution of amounts was manipulated between participants, and the distribution of delays was held constant.

6.6.1. Method. Participants. Forty Warwick psychology first-year undergraduates participated for course credit. In addition, participants knew that one participant would be drawn at random and paid according to their selection on one randomly drawn choice. Payment was by bank transfer on the day set by their choice of delay. After applying an experiment exchange rate, participants could win up to £20.

Data from four 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. Delayed options were made by crossing a set of delays with a set of amounts. All participants received delays of one day, two days, one week, two weeks, one month, two months, six months, and one year. The distribution of amounts was either positively or negatively skewed, with values from Experiment 1 (see Figure 3, center panel). One hundred twenty nondominated choices were selected at random and combined with 30 stochastically dominated choices selected at random.

6.6.2. Results and Discussion. The analysis repeats the procedure from Experiment 1, with the subjective values of amounts estimated by maximum likelihood; that is, we fitted

$$\text{Prob}(\text{sooner}) = \frac{\text{bias}_{\text{cond}} [v_{\text{cond}}(y)/u]^{\gamma_{\text{cond}}}}{\text{bias}_{\text{cond}} [v_{\text{cond}}(y)/u]^{\gamma_{\text{cond}}} + [v_{\text{cond}}(x)/t]^{\gamma_{\text{cond}}}}, \quad (5)$$

where y at time u is a smaller sooner offer, and x at time t is a later larger offer; $v_{\text{cond}}(x)$ gives the utility of amount x in condition cond . Figure 6(a) shows the revealed utility functions. The utility functions 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 concave in the positive-skew condition than the negative-skew condition ($\chi^2(1) = 4.49$, $p = 0.034$). As in Experiment 1A, £200 receives a higher valuation in the positive-skew condition than £310 in the negative-skew condition ($\chi^2(1) = 3.89$, $p = 0.047$). (The significant difference in the valuation of £200 and £310 is only marginally significant when *bias* is held constant across conditions.)

6.7. Experiment 3B

6.7.1. Method. Experiment 3B is a replication of Experiment 3A. Nineteen first-year Warwick undergraduates participated for course credit. The incentive for the randomly drawn participant was up to £50. The only other difference from Experiment 3A was that we used an even more positively skewed set of common delays (one day, two days, one week, one month, and one year). No participants violated dominance on more than 10% of catch trials, so all data were retained.

6.7.2. Results and Discussion. Figure 6(b) plots the utility functions. The results are very similar, with a significant difference in concavity ($\chi^2(1) = 6.63$, $p = 0.010$), though the conservative £200–£310 comparison only approached significance ($\chi^2(1) = 3.50$, $p = 0.061$).

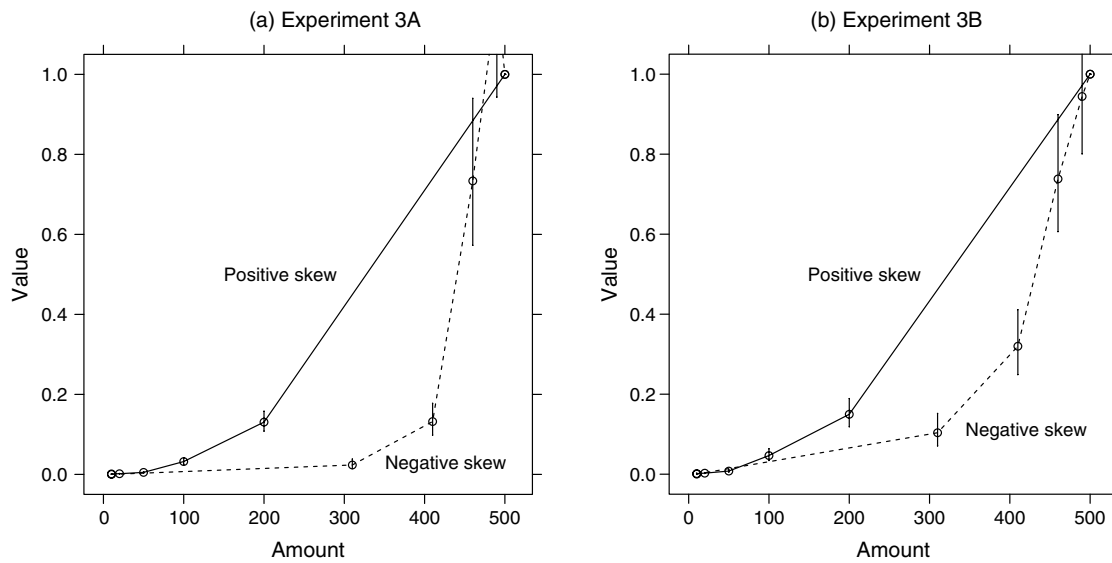
6.8. Experiment 4

6.8.1. 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 Experiments 3A and 3B.

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. Options were made by crossing a set of delays 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).

Figure 6 Utility Functions for Experiments 3A and 3B



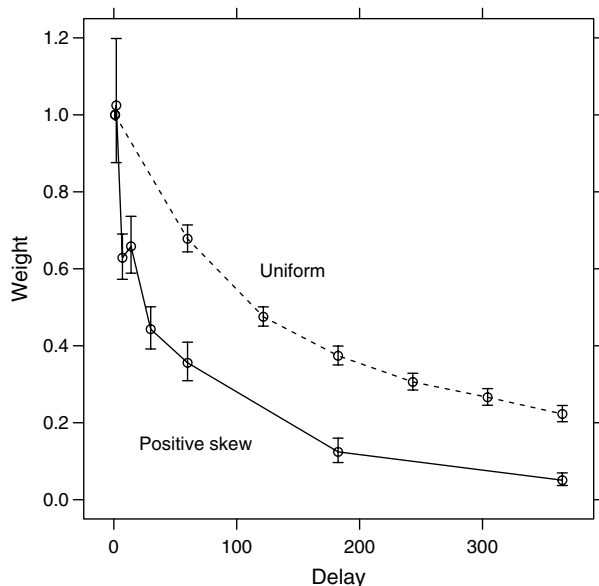
Note. Error bars are standard errors.

6.8.2. Results and Discussion. The analysis repeats the procedure from earlier experiments, with weightings of delays allowed to vary between conditions; that is, we fitted

$$\text{Prob}(\textit{sooner}) = \frac{\textit{bias}_{\textit{cond}}[w_{\textit{cond}}(u)y]^{\gamma_{\textit{cond}}}}{\textit{bias}_{\textit{cond}}[w_{\textit{cond}}(u)y]^{\gamma_{\textit{cond}}} + [w_{\textit{cond}}(t)x]^{\gamma_{\textit{cond}}}}, \quad (6)$$

where $w_{\textit{cond}}(t)$ gives the weighting of delay t in condition \textit{cond} . Figure 7 shows the revealed delay discounting functions. The delay discounting function is

Figure 7 The Delay Discounting Functions for Experiment 4



Note. Error bars are standard errors.

initially much steeper in the positive-skew condition and much closer to linear in the uniform condition. Modeling with discounting functions constrained to follow power laws showed a significant difference in curvature ($\chi^2(1) = 22.25, p < 0.0001$). Modeling with free weights revealed that, for the common two-month and six-month delays, people behaved as if they weighted delayed amounts less heavily in the positive-skew condition compared to the uniform condition, $\chi^2(1) = 15.45, p < 0.0001$.

We note here the difference between this discounting rank effect and subadditivity. Subadditivity is the empirical finding that discounting measured over a larger interval is less than that estimated by multiplying discounting over constituent subintervals (e.g., Scholten and Read 2006). Here we find more discounting of delays in conditions where the distribution is more positively skewed. So subadditivity is about whether discounting within a condition is consistent when estimated over smaller or larger intervals, whereas the present discounting rank effect is about how discounting changes across conditions. The theoretical account of subadditivity is also different. Scholten and Read (2006, 2010) explain subadditivity by assuming that delay differences rather than absolute delays are the psychological primitives that are transformed by a stable discounting function, and offer no account of these rank effects. The decision-by-sampling account of rank effects works without any such transformation (Vlaev et al. 2011).

7. General Discussion

Experimentally manipulating the distribution of money (Experiments 1 and 3), probabilities (Experiment 2), and delays (Experiment 4) people experience

alters the choices people make, which in turn alters the psychoeconomic functions constructed to describe those choices. The effect of the distribution of attribute values on the subjective function for those attribute values was the same for money, probability, and time: The subjective function was steepest when the distribution of attribute values was most dense, so that more positively skewed distributions resulted in concave functions. The shape of revealed utility, weighting, and discounting functions is, at least in part, a property of the question set and not the individual.

7.1. Psychological Implications

The literature on choices between gambles documents many choice set effects and other violations of expected utility theory. One fix is to change the psychological primitive (e.g., from wealth to gains/losses (Kahneman et al. 1991, Kahneman and Tversky 1979) or from cumulative probability distributions to risk-reward branches (Birnbaum 2008)). Another fix is to assume more complexity in the transformation of objective into subjective values (e.g., with the four-fold pattern motivating the inverse-S weighting function, Kahneman and Tversky 1979). The effects we present here are different. They are not accounted for by a change of primitive or a change in the transformation between objective and subjective values. So accounts based around the expected or discounted utility framework do not explain the results. Psychologically, the translation between objective and subjective values is not well modeled as a lookup process using some stable function or table. We have proposed the decision-by-sampling model as an explanation in which attribute values are evaluated against other attribute values in working memory. Because working memory reflects the distribution of money, probability, and delay in the environment, valuations will vary across experimental or environmental distributions.

The process claims about memory embodied in decision by sampling may offer an explanation for the link between individual differences in memory and intelligence and decision making. Higher intelligence, which is associated with higher working memory capacity, is associated with less risk aversion and less discounting (Benjamin et al. 2013, Burks et al. 2009, Dohmen et al. 2010, Shamosh and Gray 2008), and experimentally reducing working memory capacity increases discounting and risk aversion (Hinson et al. 2003, Whitney et al. 2008). As speculation, we propose that decision by sampling may predict these effects because those with lower working memory capacity will be more dominated by the immediate experimental context (which typically will induce risk aversion and discounting; see below) and less influenced by extraexperiment attribute values.

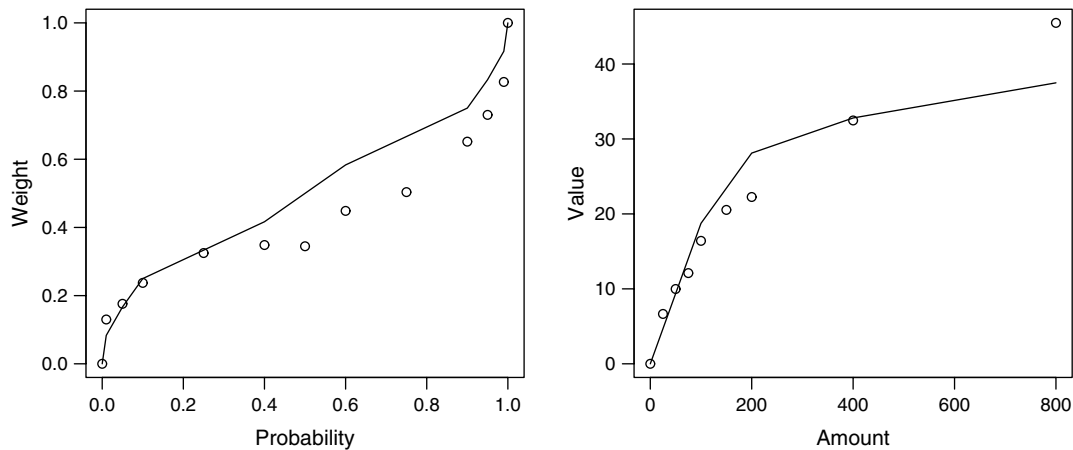
7.2. Application in Economics

Gul and Pesendorfer (2005) explain that economists are often only interested in choice behaviour, because only actual choices affect institutions and markets. Thus, evidence from psychology and neuroeconomics about brain states or about the difference between what people actually choose and what they wanted to choose is not relevant. So while the decision-by-sampling explanation may not be of direct interest to economics, our experiments—which are just about choices—should be. The experiments show that, to describe an individual, the three single curves for the utility of money and the weighting of delays and probabilities in Figure 1 need to be replaced with three books of curves with pages for each attribute-distribution context. However, using a separate function for each possible context leaves the relationship between the attribute-value distribution and the function unexplained.

The models of Kőszegi and Rabin (2006, 2007) and Maccheroni et al. (2009a, b) both incorporate the notion that the subjective value of an attribute depends on the distribution of attribute values. These models differ from decision by sampling in including a stable classical utility component and, for Kőszegi and Rabin (2006, 2007), their comparison to the reference distribution is cardinal, not ordinal. But, like decision by sampling, both of these models involve a comparison to a reference distribution: Kőszegi and Rabin (2006, 2007) assume the distribution is of (rational) expected outcomes, Maccheroni et al. (2009a, b) assume it is the outcomes of peers, and Stewart et al. (2006) assume it is the attribute values in memory—and these values could well include expected outcomes or the outcomes of others.

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

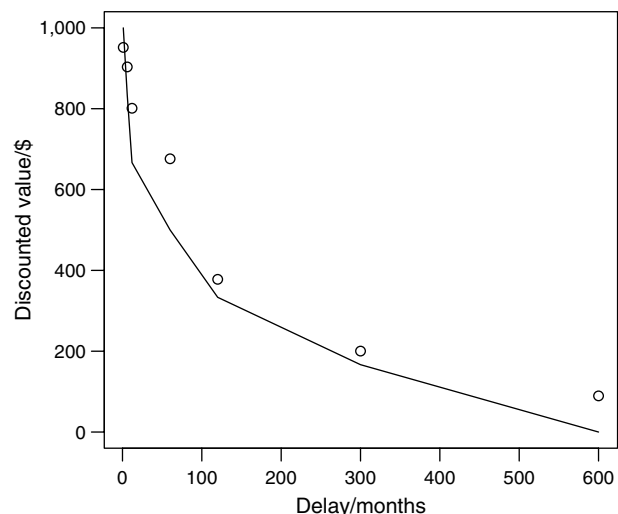
Previous investigations tend to find concave utility functions, inverse-S-shaped probability weighting functions, and hyperbolic-like discounting functions. Why? The answer, we think, is in the choice of attribute-value distributions used in these experiments. The essential observation here is that, 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. The open circles in Figure 8 plot 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

Figure 8 Empirically Derived Utility (Left) and Weighting (Right) Functions from Gonzalez and Wu (1999)

Notes. Circles plot the functions that Gonzalez and Wu recovered. The solid line plots the predictions of decision-by-sampling theory.

the utility and weighting functions, the functions predicted under this simple rank hypothesis show the same qualitative pattern seen in the empirical functions: Because there were more small amounts than large amounts, the utility function is steeper for small amounts and flatter for large amounts, and because there were more small and large probabilities than intermediate probabilities, the weighting function is steeper for small and large probabilities than for intermediate probabilities. Because Gonzalez and Wu (1999) used a distribution of probabilities symmetrical around 0.5, the rank hypothesis does not predict the asymmetry in the empirical weighting function. But, if participants also recalled probabilities from outside the experiment from a distribution with more larger probabilities, like that in Figure 2, the rank hypothesis does predict the asymmetry. Therefore, there is a self-fulfilling result here: By taking more measurements where previous research indicates functions changed most quickly, this will lead to steeper functions where more measurements are taken.

The same logic can be applied to classic studies in intertemporal choice. Rachlin et al. (1991) used hypothetical choices between an immediately available sum of money and a delayed \$1,000. The open circles in Figure 9 plot the discounted value as a function of delay. The line on the plot shows the function predicted by decision by sampling assuming that the set of delays presented provides the comparison set. As for the Gonzalez and Wu (1999) data, the function predicted under this simple rank hypothesis shows the same qualitative pattern seen in the empirical function. The discounting function is roughly hyperbolic because of the geometrically spaced distribution of delays chosen by Rachlin et al. (1991). Virtually all discounting studies use this kind of distribution of delay.

Figure 9 An Empirically Derived Discounting Function from Rachlin et al. (1991), Showing the Discounted Value of a Delayed \$1,000

Notes. Circles plot the function that Rachlin et al. recovered. The solid line plots the predictions of decision-by-sampling theory.

7.4. Neuroeconomic Evidence

Our experimental results are consistent with recent evidence of coding of relative value within the brain (for a review, see Seymour and McClure 2008). 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 reward available on that choice. In contrast, Padoa-Schioppa and Assad (2008) find choice-invariant responding.

With intermixed choices between quantities of three juices, macaque orbitofrontal cortex neurons respond to the absolute quantity of juice available, irrespective of which other juice is presented in the choice. The value of a juice was coded independently of the other juice available on that choice. These two findings are reconciled by noting that Tremblay and Schultz (1999) blocked their choice pairs, but choice pairs were randomly intermixed for Padoa-Schioppa and Assad (2008). Therefore, the findings of relative value responding by Tremblay and Schultz (1999) and choice-invariant responding by Padoa-Schioppa and Assad (2008) are consistent with orbitofrontal cortex neurons encoding the value of rewards relative to other rewards within the current block, rather than just within the immediate choice.

Mullett and Tunney (2013) find exactly this block dependency using functional magnetic resonance imaging in humans. In some blocks, participants saw prizes of £0.10, £0.20, and £0.30, and in other blocks participants saw prizes of £5.00, £7.00, and £10.00. The activity in the ventral striatum and thalamus was sensitive to the relative value of a prize within a block. The activity in the ventral medial prefrontal cortex and the anterior cingulate cortex was sensitive to the rank of the prize across the whole experiment. Mullett and Tunney (2013) can discriminate between sensitivity to value and sensitivity to rank because across the whole experiment the distribution of prizes is not uniform. As Mullett and Tunney (2013) conclude, this pattern is as predicted by decision by sampling: prizes from the experiment enter memory and form the context against which a current prize is judged, with representations proportional to the rank position of the prize within the experiment context rather than the absolute value of the prize.

7.5. Conclusion

Manipulation of the distributions of gains, risks, and delays people experience systematically changes the utility, weighting, and discounting functions revealed from people’s choices. So, if we can measure lots of utility, weighting, or discounting functions just by changing the choices used to estimate them (whether in the lab or the field), which functions are the “true” functions? It may be that there is no stability within a person and that some process like decision by sampling is all there is, with the apparent stability coming only from stability of the distribution of attribute values over time. In this case, there are no “true” functions. Or it may be that people do have some stable underlying mapping between objective and subjective values. This is an empirical matter, but, at the least, this mapping must be quite malleable: In our data the effect of the attribute-value distributions is about

as large as the individual differences between people. Decision by sampling provides a common account of the origin of utility, weighting, and discounting functions—and the explanation of how we were able to change their shapes.

Supplemental Material

Supplemental material to this paper is available at <http://dx.doi.org/10.1287/mnsc.2013.1853>.

Acknowledgments

The authors thank James Adelman for valuable help with the logistic regression analysis and Gordon D. A. Brown, Emina Canic, Tatiana Kornienko, Graham Loomes, Will Matthews, Timothy L. Mullett, Dani Navarro-Martinez, Takao Noguchi, Ganna Pogrebna, and Christoph Ungemach for insightful discussion. This research was supported by the Economic and Social Research Council [Grants RES-062-23-0952, RES-000-22-3339, ES/K002201/1]; and the Leverhulme Trust [Grant RP2012-V-022].

Appendix A. Estimation of Utilities and Weights

In Experiments 1A–1C, we fitted Equation (A1) to estimate the utilities of each amount:

$$\log \left[\frac{\text{Prob}(\text{safe})}{1 - \text{Prob}(\text{safe})} \right] = \nu + \tau \text{cond} + \omega \log \left(\frac{q}{p} \right) + \xi \text{cond} \log \left(\frac{q}{p} \right) + \sum_i \beta_i X_i. \quad (\text{A1})$$

The first term, ν , is an intercept. In the second term, τ is the coefficient for a dummy variable cond ; cond indicates condition, with 0 for positive skew and 1 for negative skew or uniform. In the third term, ω is the coefficient for $\log(q/p)$, the log of the ratio of the safe and risky probabilities. In the fourth term, ξ is the coefficient for the interaction of cond and $\log(q/p)$. In the last term, the X_i ’s are dummy variables indicating the presence of each amount_i ; X_i is +1 if amount_i appears in the safe gamble, –1 if amount_i appears in the risky gamble, or 0 if amount_i does not appear in the choice. The β_i ’s are coefficients for the amount dummies.

Though it is perhaps not obvious, Equation (A1) is a rearrangement of Equation (3) if we set

$$\log(\text{bias}_{\text{cond}}) = \nu + \text{cond} \tau, \quad (\text{A2})$$

$$\gamma_{\text{cond}} = \omega + \text{cond} \xi, \quad (\text{A3})$$

and

$$u_{\text{cond}}(\text{amount}_i) = \exp(\beta_i / \gamma_{\text{cond}}). \quad (\text{A4})$$

The advantage of Equation (A1) is that this is the standard form for a logistic regression with the choice on each trial as the dependent variables and cond , $\log(q/p)$, and the X_i ’s as independent variables. If one treats the data as coming from one single participant and fits Equation (3) by maximum likelihood, exactly the same parameter estimates are obtained as when Equation (A1) is fitted as a logistic regression within any standard statistics package.

Of course, it is not appropriate to ignore the within-subjects nature of the design. Responses will be correlated

within each participant, and this means that responses are not independent. But because we can use the logistic regression form in Equation (A1), we are able to use off-the-shelf methods to deal with the repeated observations for each participant by fitting the model as a mixed effect logistic regression. The fixed effects part of the model is given by Equation (A1). The random effects part of the model included random intercepts and full random slopes for each participant. It is the $u_{cond}(amount_i)$ values that are of primary concern because these indicate the utilities of the amounts in each condition. These are the estimates plotted in Figure 4.

For Experiments 2A and 2B, where probability weights are estimated, amounts are swapped with probabilities so that the roles of p and x and of q and y are exchanged in Equation (A1) to give

$$\log\left[\frac{\text{Prob}(safe)}{1 - \text{Prob}(safe)}\right] = \nu + \tau \text{cond} + \omega \log\left(\frac{y}{x}\right) + \xi \text{cond} \log\left(\frac{y}{x}\right) + \sum_i \beta_i X_i \quad (\text{A5})$$

where the X_i 's are now dummies for the probabilities and $w_{cond}(p) = \exp(\beta_i/\gamma_{cond})$.

For Experiments 3A and 3B, where utilities are estimated from intertemporal choices, p is exchanged with $1/t$, where t is the longer delay, and q is exchanged with $1/u$ where, u is the shorter delay, to give

$$\log\left[\frac{\text{Prob}(safe)}{1 - \text{Prob}(safe)}\right] = \nu + \tau \text{cond} + \omega \log\left(\frac{t}{u}\right) + \xi \text{cond} \log\left(\frac{t}{u}\right) + \sum_i \beta_i X_i. \quad (\text{A6})$$

For Experiment 4, where delay weights are estimated, we used

$$\log\left[\frac{\text{Prob}(safe)}{1 - \text{Prob}(safe)}\right] = \nu + \tau \text{cond} + \omega \log\left(\frac{y}{x}\right) + \xi \text{cond} \log\left(\frac{y}{x}\right) + \sum_i \beta_i X_i. \quad (\text{A7})$$

Appendix B. Using Parametric Forms to Test Curvature

In Experiments 1A–1C, we fitted Equation (B1) to test for differences in the curvature of the utility functions:

$$\log\left[\frac{\text{Prob}(safe)}{1 - \text{Prob}(safe)}\right] = \nu + \tau \text{cond} + \omega \log\left(\frac{q}{p}\right) + \xi \text{cond} \log\left(\frac{q}{p}\right) + \zeta \log\left(\frac{y}{x}\right) + \eta \text{cond} \log\left(\frac{y}{x}\right). \quad (\text{B1})$$

Here we set

$$\log(bias_{cond}) = \nu + \text{cond} \tau, \quad (\text{B2})$$

$$\gamma_{cond} = \omega + \text{cond} \xi, \quad (\text{B3})$$

and

$$\alpha_{cond} = \frac{\zeta + \text{cond} \eta}{\gamma_{cond}}. \quad (\text{B4})$$

If the unspecified utility function in Equation (3) is replaced with a power function $u_{cond}(x) = x^{\alpha_{cond}}$, then Equation (B1) is just a rearrangement of Equation (3), and, as above, fitting Equations (3) and (B1) is equivalent.

As above, we deal with the repeated observations of each participant by fitting the model as a mixed effects logistic regression. The fixed effects part of the model is given by Equation (B1). The random effects part of the model included random intercepts and full random slopes for each participant. It is the estimates of α_0 and α_1 and the difference between them that gives the critical test of the difference in utility function curvature between conditions.

For Experiments 2A and 2B, we fit

$$\log\left[\frac{\text{Prob}(safe)}{1 - \text{Prob}(safe)}\right] = \nu + \tau \text{cond} + \omega \log\left(\frac{y}{x}\right) + \xi \text{cond} \log\left(\frac{y}{x}\right) + \zeta \log\left(\frac{q}{p}\right) + \eta \text{cond} \log\left(\frac{q}{p}\right), \quad (\text{B5})$$

where $w_{cond}(p) = p^{\alpha_{cond}}$, and $bias$, γ , and α are given by Equations (B2)–(B4).

For Experiments 3A and 3B, we fit

$$\log\left[\frac{\text{Prob}(safe)}{1 - \text{Prob}(safe)}\right] = \nu + \tau \text{cond} + \omega \log\left(\frac{t}{u}\right) + \xi \text{cond} \log\left(\frac{t}{u}\right) + \zeta \log\left(\frac{y}{x}\right) + \eta \text{cond} \log\left(\frac{y}{x}\right), \quad (\text{B6})$$

where $u_{cond}(x) = x^{\alpha_{cond}}$, and $bias$, γ , and α are given by Equations (B2)–(B4). For Experiment 4, we fit

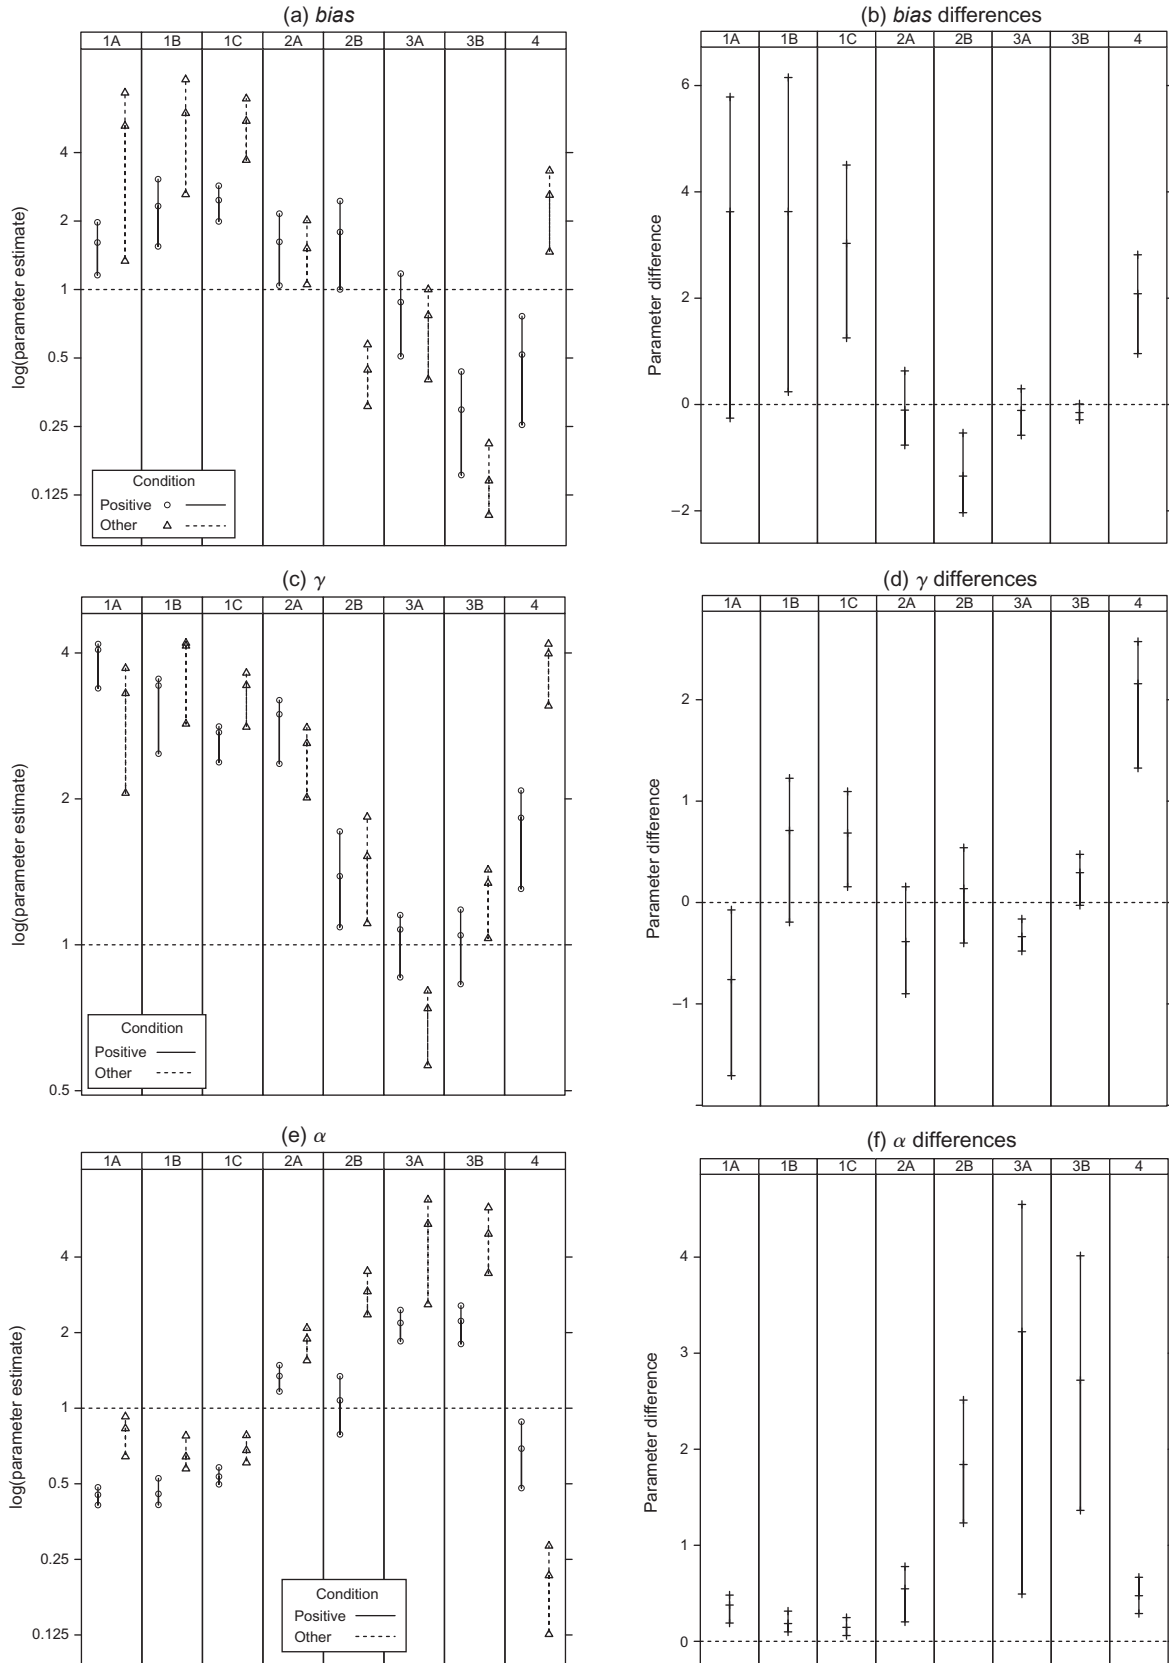
$$\log\left[\frac{\text{Prob}(safe)}{1 - \text{Prob}(safe)}\right] = \nu + \tau \text{cond} + \omega \log\left(\frac{y}{x}\right) + \xi \text{cond} \log\left(\frac{y}{x}\right) + \zeta \log\left(\frac{t}{u}\right) + \eta \text{cond} \log\left(\frac{t}{u}\right), \quad (\text{B7})$$

where $w_{cond}(t) = 1/t^{\alpha_{cond}}$, and $bias$, γ , and α are given by Equations (B2)–(B4).

For each experiment, Figure B.1 shows the estimates of $bias$, γ , and α for each condition together with bootstrapped 95% confidence intervals. The left column shows the parameter estimates with a solid confidence interval for the positive condition and a dashed confidence interval for the other condition (uniform or negative, depending on the experiment). There are panels for each experiment. The value 1 is marked with a horizontal dotted line. For $bias$, 1 indicates no bias; for γ , 1 indicates probability matching; and for α , 1 indicates a linear utility function, linear probability weighting function, or hyperbolic delay discounting. The right column shows a confidence interval for the difference in the parameter estimates. The top row plots $bias$ and $bias$ differences. That the red confidence intervals do not include zero (the dotted line) for $bias$ difference estimates for Experiments 1B, 1C, 2B, and 4 indicates significant differences in $bias$ in these studies. The middle row shows γ and γ differences. Of most interest is the bottom row, which shows α and α differences. The 95% confidence intervals for α differences are all above zero, which indicates that in every experiment there is a significant difference in the curvature of the utility, weighting, or discounting function.

Raw data and full R source code are available from the authors.

Figure B.1 Estimates and 95% Confidence Intervals for Parameter Estimates (Left) and the Difference Between Parameters Between Conditions (Right)



Downloaded from informs.org by [137.205.50.42] on 09 April 2015, at 07:46. For personal use only, all rights reserved.

References

- Ainslie G (1975) Specious reward: A behavioural theory of impulsiveness and impulse control. *Psychol. Bull.* 82(5):463–496.
- Allais M (1953) Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école américaine [Rational man's behavior in face of risk: Critique of the American school's postulates and axioms]. *Econometrica* 21(4): 503–546.
- Beauchamp JP, Benjamin DJ, Chabris CF, Laibson DI (2012) How malleable are risk preferences and loss aversion? Working paper, Harvard University, Cambridge, MA. http://economics.cornell.edu/dbenjamin/HowMalleableareRiskPreferencesandLossAversion_3-10-2012-Combined.pdf.
- Benartzi S, Thaler RH (2001) Naive diversification strategies in defined contribution saving plans. *Amer. Econom. Rev.* 91(1): 79–98.
- Benjamin DJ, Brown SA, Shapiro JM (2013) Who is “behavioral”? Cognitive ability and anomalous preferences. *J. Eur. Econom. Assoc.* 11(6):1231–1255.
- Birnbaum MH (1992) Violations of the monotonicity and contextual effects in choice-based certainty equivalents. *Psychol. Sci.* 3(5): 310–314.
- Birnbaum MH (2008) New paradoxes of risky decision making. *Psychol. Rev.* 115(2):453–501.
- Birnbaum MH, Chavez A (1997) Tests of theories of decision making: Violations of branch independence and distribution independence. *Organ. Behav. Human Decision Process.* 71(2):161–194.
- Boyce CJ (2009) Subjective well-being: An intersection between economics and psychology. Ph.D. thesis, Warwick University, Coventry, UK.
- Brown GDA, Gardner J, Oswald AJ, Qian J (2008) Does wage rank affect employees' well-being? *Indust. Relat.* 47(3):355–389.
- Burks SV, Carpenter JP, Goette L, Rustichini A (2009) Cognitive skills affect economic preferences, strategic behavior, and job attachment. *Proc. Natl. Acad. Sci. USA* 106(19):7745–7750.
- Busemeyer JR, Townsend JT (1993) Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychol. Rev.* 100(3):432–459.
- Camerer CF (1995) Individual decision making. Kagel J, Roth AE, eds. *Handbook of Experimental Economics* (Princeton University Press, Princeton, NJ), 587–703.
- Dohmen T, Falk A, Huffman D, Sunde U (2010) Are risk aversion and impatience related to cognitive ability? *Amer. Econom. Rev.* 100(3):1238–1260.
- Edwards W (1962) Subjective probabilities inferred from decisions. *Psychol. Rev.* 69(2):109–135.
- Gonzalez R, Wu G (1999) On the shape of the probability weighting function. *Cognitive Psychol.* 38(1):129–166.
- Gul F, Pesendorfer W (2005) The case for mindless economics. Levine's working paper archive, Princeton University, Princeton, NJ. <http://www.dklevine.com/archive/refs478482800000000581.pdf>.
- Haggag K, Paci G (2013) Default tips. Working paper, Columbia University, New York. <http://hdl.handle.net/10022/AC:P:19386>.
- Hinson JM, Jameson TL, Whitney P (2003) Impulsive decision making and working memory. *J. Experiment. Psychol.: Learn.* 29(2): 298–306.
- Kahneman D, Tversky A (1979) Prospect theory: An analysis of decision under risk. *Econometrica* 47(2):263–291.
- Kahneman D, Knetsch JL, Thaler RH (1991) Anomalies: The endowment effect, loss aversion, and status quo bias. *J. Econom. Perspect.* 5(1):193–206.
- Kornienko T (2011) A cognitive basis for context-dependent utility. Working paper, University of Edinburgh, Edinburgh, UK. doi: 10.1.1.296.7140.
- Kőszegi B, Rabin M (2006) A model of reference-dependent preferences. *Quart. J. Econom.* 121(4):1133–1165.
- Kőszegi B, Rabin M (2007) Reference-dependent risk attitudes. *Amer. Econom. Rev.* 97(4):1047–1073.
- Loewenstein G, Prelec D (1992) Anomalies in intertemporal choice: Evidence and an interpretation. *Quart. J. Econom.* 107(2): 573–597.
- Loomes G (2010) Modeling choice and valuation in decision experiments. *Psychol. Rev.* 117(3):902–924.
- Loomes G, Sugden R (1982) Regret theory: An alternative theory of rational choice under uncertainty. *Econom. J.* 92(368):805–824.
- Luce RD (1959) *Individual Choice Behavior* (John Wiley & Sons, New York).
- Luce RD (2000) *Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches* (Erlbaum, Mahwah, NJ).
- Maccheroni F, Marinacci M, Rustichini A (2009a) Pride and diversity in social economics. Working paper, University of Bocconi, Milan, Italy. http://www.igier.unibocconi.it/files/pride_and_Diversity.pdf.
- Maccheroni F, Marinacci M, Rustichini A (2009b) Social decision theory: Choosing within and between groups. Working paper, University of Bocconi, Milan, Italy. <http://ideas.repec.org/p/cca/wpaper/71.html>.
- Matthews WJ (2012) How much do incidental values affect the judgment of time? *Psychol. Sci.* 23(11):1432–1434.
- Matthews WJ (2013) Personal communication with Neil Stewart, June 20–21.
- Mazur JE (1987) An adjusting procedure for studying delayed reinforcement. Commons ML, Mazur JE, Nevin JA, Rachlin H, eds. *Quantitative Analyses of Behavior: The Effect of Delay and of Intervening Events on Reinforcement Value*, Vol. 5 (Erlbaum, Hillsdale, NJ), 55–73.
- Mazar N, Kőszegi B, Ariely D (2014) True context-dependent preferences? The causes of market-dependent valuations. *J. Behav. Decision Making* 27(3):200–208.
- Mullett TL (2012) Personal communication with Neil Stewart, December 17–18.
- Mullett TL, Tunney RJ (2013) Value representations by rank order in a distributed network of varying context dependency. *Brain Cognition* 82(1):76–83.
- Padoa-Schioppa C, Assad JA (2008) The representation of economic value in the orbitofrontal cortex is invariant for changes of menu. *Nat. Neurosci.* 11(1):95–102.
- Parducci A (1965) Category judgment: A range-frequency model. *Psychol. Rev.* 72(6):407–418.
- Parducci A (1995) *Happiness, Pleasure and Judgment: The Contextual Theory and Its Applications* (Erlbaum, Mahwah, NJ).
- Peetz J, Wilson AE (2013) The post-birthday world: Consequences of temporal landmarks for temporal self-appraisal and motivation. *J. Personality Soc. Psychol.* 104(2):249–266.
- Quiggin J (1993) *Generalized Expected Utility Theory: The Rank-Dependent Model* (Kluwer Academic Publishers, Norwell, MA).
- Rachlin H, Raineri A, Cross D (1991) Subjective probability and delay. *J. Experiment. Anal. Behav.* 55(2):233–244.
- Samuelson PA (1937) A note on measurement of utility. *Rev. Econom. Stud.* 4(2):155–161.
- Savage LJ (1954) *The Foundations of Statistics* (John Wiley & Sons, New York).
- Schoemaker PJH (1982) The expected utility model: Its variants, purposes, evidence, and limitations. *J. Econom. Literature* 20(2):529–563.
- Scholten M, Read D (2006) Discounting by intervals: A generalized model of intertemporal choice. *Management Sci.* 52(9): 1424–1436.
- Scholten M, Read D (2010) The psychology of intertemporal trade-offs. *Psychol. Rev.* 117(3):925–944.
- Seymour B, McClure SM (2008) Anchors, scales and the relative coding of value in the brain. *Curr. Opin. Neurobiol.* 18(2): 173–178.

- Shamosh NA, Gray JR (2008) Delay discounting and intelligence: A meta-analysis. *Intelligence* 36(4):289–305.
- Shepard RN (1957) Stimulus and response generalization: A stochastic model relating generalization to distance in psychological space. *Psychometrika* 22(4):325–345.
- Starmer C (2000) Developments in non-expected utility theory: The hunt for a descriptive theory of choice under risk. *J. Econom. Lit.* 38(2):332–382.
- Stewart N (2009) Decision by sampling: The role of the decision environment in risky choice. *Quart. J. Experiment. Psychol.* 62(6):1041–1062.
- Stewart N, Simpson K (2008) A decision-by-sampling account of decision under risk. Chater N, Oaksford M, eds. *The Probabilistic Mind: Prospects for Bayesian Cognitive Science* (Oxford University Press, Oxford, UK), 261–276.
- Stewart N, Chater N, Brown GDA (2006) Decision by sampling. *Cognitive Psychol.* 53(1):1–26.
- Stewart N, Chater N, Stott HP, Reimers S (2003) Prospect relativity: How choice options influence decision under risk. *J. Experiment. Psychol.: General* 132(1):23–46.
- Thaler RH (1981) Some empirical evidence on dynamic inconsistency. *Econom. Lett.* 8(3):201–207.
- Tremblay L, Schultz W (1999) Relative reward preference in primate orbitofrontal cortex. *Nature* 398(6729):704–708.
- Tversky A, Kahneman D (1992) Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertainty* 5(4):297–323.
- Ungemach C, Stewart N, Reimers S (2011) How incidental values from our environment affect decisions about money, risk, and delay. *Psychol. Sci.* 22(2):253–260.
- Vlaev I, Chater N, Stewart N, Brown GDA (2011) Does the brain calculate value? *Trends Cogn. Sci.* 15(11):546–554.
- Vlaev I, Seymour B, Dolan RJ, Chater N (2009) The price of pain and the value of suffering. *Psychol. Sci.* 20(3):309–317.
- von Neumann M, Morgenstern O (1947) *Theory of Games and Economic Behavior*, 2nd ed. (Princeton University Press, Princeton, NJ).
- Whitney P, Rinehart CA, Hinson JM (2008) Framing effects under cognitive load: The role of working memory in risky decisions. *Psychon. B. Rev.* 15(6):1179–1184.



This work is licensed under a Creative Commons Attribution 3.0 United States License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “Management Science. Copyright 2015 INFORMS. <http://dx.doi.org/10.1287/mnsc.2013.1853>, used under a Creative Commons Attribution License: <http://creativecommons.org/licenses/by/3.0/us/>.”