THE UNIVERSITY OF
WARWICK

University of Warwick institutional repository: http://go.warwick.ac.uk/wrap

This paper is made available online in accordance with publisher policies. Please scroll down to view the document itself. Please refer to the repository record for this item and our policy information available from the repository home page for further information.

To see the final version of this paper please visit the publisher's website. Access to the published version may require a subscription.

Decision by Sampling: The Role of the Decision Environment in Risky Choice

Neil Stewart
University of Warwick

Decision by sampling (DbS) is a theory about how our environment shapes the decisions that we make. Here, I review the application of DbS to risky decision making. According to classical theories of risky decision making, people make stable transformations between outcomes and probabilities and their subjective counterparts using fixed psychoeconomic functions. DbS offers a quite different account. In DbS, the subjective value of an outcome or probability is derived from a series of binary, ordinal comparisons with a sample of other outcomes or probabilities from the decision environment. In this way, the distribution of attribute values in the environment determines the subjective valuations of outcomes and probabilities. I show how DbS interacts with the real-world distributions of gains, losses, and probabilities to produce the classical psychoeconomic functions. I extend DbS to account for preferences in benchmark data sets. Finally, in a challenge to the classical notion of stable subjective valuations, I review evidence that manipulating the distribution of attribute values in the environment changes our subjective valuations just as DbS predicts.

Risky decision making is a central part of human cognition. One often has to choose between alternative actions where the outcomes associated with each action are uncertain. Because our environment is not a deterministic place, most everyday decisions involve some element of risk. And many of our most important decisions also involve risk. For example, financial decisions involving saving and borrowing are risky because of variability in interest rates and the stock market. Medical decisions are risky because the effectiveness of treatments will vary from case to case.

In this review, I show how the risky decisions that we make are influenced by the statistical distributions of risks and rewards in the environment. Our sensitivity to the distribution of attributes within the environment emerges from our use of a set of domain-general cognitive tools to make risky decisions and is captured in the decision by sampling model (DbS, Stewart, Chater, & Brown, 2006). I begin by reviewing what might be considered to be the most prominent theory of risky decision making: prospect theory (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992). Prospect theory provides an excellent description of human risky decision making. In the second part of the article, I show how it is possible to derive the prospect-theory description from the real-world distribution of risks and rewards under the assumptions of the DbS model. In the penultimate part of this article, I show how the DbS model has been successfully developed to provide a process account of the risky decisions that we make. Finally, in a direct test of the DbS model, I review new experimental evidence for the link between the distribution of attribute values in the environment and the risky decisions that we make.

Descriptive Models of Risky Decision Making

If you were offered a choice between either (a) £1,000 or (b) a 50% chance of £2,000 otherwise £0 which would you choose? Questions like these have been used extensively in the study of human risky decision making as carefully controlled proxies for real-life risky decisions. Lópes and Oden (1983), and later Kahneman (Kahneman & Tversky, 2000), make the analogy between the study of choices between simple gambles in risky decision making and the study of the fruit fly in genetics. Returning to the example, the majority of people are risk averse and have a preference for the certain £1,000. This result is immediately useful in ruling out perhaps the most obvious candidate—expected value—as a model for human risky decision making.

In expected value theory, people prefer the option which offers them the highest average payoff. For the first option, the expected value is £1,000. For the second option, the expected values is also £1,000 = $^1/_2 \times £2,000 + ^1/_2 \times £0$. The two options have the same expected value and so expected value does not capture the preference for the first option.

Economists have a theory that does capture this risk-averse preference: expected utility (EU) theory (Bernoulli,
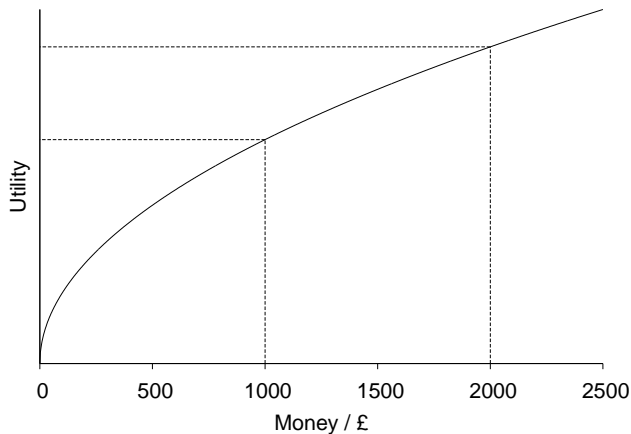
Figure 1. A power law utility function $U(x) = x^\alpha$ transforms money x into its subjective equivalent $U(x)$.

1738/1954; see von Neumann & Morgenstern, 1947 for an axiomatisation). In EU theory, money is transformed into *utility* before expectations are taken. Figure 1 shows a typical power-law utility function where $U(x) = x^\alpha$. (Bernoulli suggested a logarithmic function, but credits Cramer, 1728, with suggesting the square root function. For investigations into functional forms for utility functions, see Bell & Fishburn, 1999; Daniels & Keller, 1992; Fishburn & Kochenberger, 1979.) The utility function has a concave-downward shape (when $\alpha < 1$ the function is curved downwards) and this allows the model to capture the risk aversion displayed in the preference for the certain £1,000. Specifically, because the utility of £1,000 is more than half the utility of £2,000 when the utility curve is concave, EU theory predicts a higher utility for, and thus a preference for, the first option. Specifically, $U(\text{£}1,000) > {}^1\!/_2 \times U(\text{£}2,000) + {}^1\!/_2 \times U(\text{£}0)$ when $\alpha < 1$.

The concave shape of the utility function captures and explains risk aversion. But there is evidence that utility functions have this concave-downward form when they are measured in risk free scenarios. For example, Galanter (1962) found a concave shape by asking participants to judge how much money would make them twice as happy as a reference amount. Other theories place risk aversion elsewhere (see Davies & Satchell, 2007, for a recent discussion).

EU theory plays a foundational role in economics, because it is used as a model of the rational individual in many models of the economy. Its mathematical simplicity has given it great appeal. Unfortunately, although it captures risk aversion, it is not a complete description of human risky decision making. There are many examples of preference patterns that violate EU theory (for reviews see Allais, 1953; Birnbaum, 2008; Camerer, 1995; Luce, 2000; Schoemaker, 1982; Starmer, 2000). Committing a grave injustice to the volume of empirical work demonstrating violations of EU theory, I will present here only two examples taken from Kahneman and Tversky (1979). First, consider the pair of choices in Table 1. These choices are an example of the common-ratio effect (Allais, 1953). Choice 1 offers either

(a) an 80% chance of £4,000 otherwise £0 or (b) £3,000. 80% of Kahneman and Tversky's participants preferred the sure £3,000. Choice 2 offers either (a) a 20% chance of £4000 otherwise £0 or (b) a 25% chance of £3,000 otherwise nothing. Now the majority of participants preferred the 20% chance of £4,000.

The preference for Prospect B in Choice 1 but Prospect A in Choice 2 represents a violation of EU theory because Choice 2 is generated from Choice 1 by multiplying probabilities by 1/4. Because the same thing was done to both prospects in the choice, preference should not switch from one side to the other. More specifically, the preference for Prospect B in Choice 1 implies $.80 \times U(\text{£}4,000) < 1.00 \times U(\text{£}3,000)$ but the preference for Prospect A in Choice 2 implies $.20 \times U(\text{£}4,000) > .25 \times U(\text{£}3,000)$. These two inequalities contradict one another and cannot both be true, and thus this pattern of preference violates EU theory. Although this violation is demonstrated at the population level, a large proportion of individual participants show the preference for Prospect B in Choice 1 and Prospect A in Choice 2 when tested with both choices (e.g., Carlin, 1992).

Kahneman and Tversky (1979) developed prospect theory to account for violations like the common-ratio effect. They deliberately kept their formulation close to EU theory, and incorporated the minimum modifications necessary to account for the data. The first key difference is the inclusion of a probability-weighting function. The left panel of Figure 2 shows the decision weighting function from Tversky and Kahneman's (1992) cumulative prospect theory. Just as real world outcomes are transformed into their subjective utility equivalents by the utility function, Kahneman and Tversky suggested that objective probabilities are transformed into subjective probabilities. The probability-weighting function is most sensitive (i.e., steepest) near 0 and 1 and is least sensitive (i.e., flattest) for mid-range probabilities. [1]

The probability-weighting function allows prospect theory to account for the common-ratio effect. Specifically, the difference between the weights of the .8 and 1.0 probabilities in Choice 1 is large, because the probability-weighting function is most sensitive in this region. But the difference between the weights of the .2 and .25 probabilities in Choice 2 is small, because the probability-weighting function is less sensitive in this region. Thus although the ratio of the objective probabilities is the same in Choices 1 and 2, the ratio of the weights is not. As a result, relatively more weight is placed on the £3,000 in Choice 1 but relatively more weight is placed on the £4,000 in Choice 2. The reversal in preference between the two choices results from the shift in the relative sizes of the decision weights.

---

[1] In fact, Kahneman and Tversky maintained a difference between subjective probability and decision weighting. A subjective probability is the psychophysical transform of the objective probability. A decision weight is the emphasis given to the corresponding outcome. So one might have an accurately calibrated subjective probability for an objectively unlikely event (e.g., knowing winning the lottery is very unlikely), but still behave as if the event is more likely than it really is because one weights the associated outcome too heavily (e.g., by buying a ticket).

Table 1
The Common-Ratio Effect

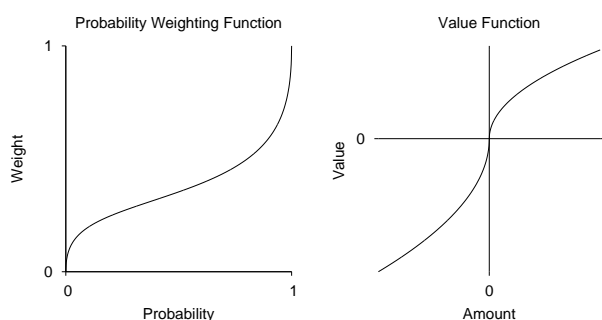| Choice	Prospect A			Prospect B		
	Amount	Probability	% chosen	Amount	Probability	% chosen
1	£4,000	.8	20%	£3,000	1.0	80%
	£0	.2				
2	£4,000	.2	65%	£3,000	.25	35%
	£0	.8		£0	.75	

Figure 2. A: Example probability-weighting function (top) and value function (bottom) from cumulative prospect theory.

A second violation of EU theory is shown in Table 2. Choice 1 pays participants an initial payment of £1,000. Participants are then offered (a) a 50% chance of winning another £1,000 otherwise £0 or (b) £500. 84% of participants prefer Prospect B. Choice 2 pays participants an initial payment of £2,000. Participants are then offered (a) a 50% chance of losing £1,000 or otherwise losing nothing or (b) a sure loss of £500. 69% of participants prefer Prospect A.

The preference for Prospect B in Choice 1 but Prospect A in Choice 2 represents a violation of EU theory because Choices 1 and 2 lead to the same net outcomes. For both choices, when initial endowments are integrated with the prospect payoffs, Prospect A offers a 50% chance of £1,000 and a 50% chance of £2,000 and Prospect B offers £1,500 for sure. Kahneman and Tversky (1979) refer to this effect as a framing effect, because whether outcomes were framed as gains or losses switched people's preferences from risk averse to risk prone.

Kahneman and Tversky (1979) incorporated this finding into prospect theory by again making a small change to the EU framework. The right panel in Figure 2 shows a revised function for transforming amounts into their subjective equivalents. In prospect theory, this function is called the value function (cf. the utility function in EU theory). In the top-right quadrant, the function is concave, just like the utility function from EU theory, to capture risk-averse behaviour in the domain of gains. In the bottom-left quadrant, the function is convex to capture risk seeking behaviour in the domain of losses. The function is also steeper for losses than for gains to capture the fact that people do not like to play zero expected value gambles involving gains and losses (but see Ert & Erev, 2007). For example, most people do not want to play a gamble that offers equal chances to gain £1,000 and to lose £1,000.

Prospect theory provides a good description of the risky decisions that people make. The shapes of the probability-weighting function and the value function are chosen to provide this good description. In addition to prospect theory, there are many other theories that have been derived from EU theory including subjective EU theory (Edwards, 1962; Savage, 1954), regret theory (Loomes & Sugden, 1982), rank-dependent utility theory (Quiggin, 1993), decision field theory (Busemeyer & Townsend, 1993), and the transfer-of-attention-exchange model (Birnbaum & Chavez, 1997; Birnbaum, 2008). What each of these theories has in common is the assumption that outcomes and probabilities are transformed into their subjective equivalents and expectations are taken. The different theories assume different transforms of amounts and probabilities, and these different transforms allow the models to describe the decision people make. In the remainder of this article I want to consider DbS. According to DbS, psychoeconomic functions have no psychological status (i.e., we do not have look-up functions inside our heads for converting between money and utility or probability and subjective probability). Instead, psychoeconomic functions are revealed from choice data: they describe the choices people make, but not the psychology of choosing.

How the Distribution of Attribute Values Shapes Revealed Psychoeconomic Functions

DbS assumes that three simple cognitive tools are the basis for decision making: binary, ordinal comparison; sampling; and frequency accumulation. Stewart et al. (2006) review the evidence for the ubiquity of these domain-general cognitive tools. Very briefly, the binary, ordinal comparison tool is motivated by findings in psychophysics, where people are rather good at saying which stimulus in a pair is the larger stimulus, but are rather bad at estimating the magnitudes of individual stimuli (Laming, 1997; Stewart, Brown, & Chater, 2005). The sampling tool is motivated by the judgement and decision making literature, where hypotheses (in norm theory, Kahneman & Miller, 1986) or uncertainties (in support theory, Tversky & Koehler, 1994) are compared to a small sample of exemplars from memory. More generally, that our working memories can hold a small sample of information from the immediate context and from long-term memory is well established. The frequency accumulation tool is motivated by the finding that we are rather good at keeping track

Table 2
The Framing Effect

Choice	Prospect A			Prospect B		
	Amount	Probability	% chosen	Amount	Probability	% chosen
1	Given £1,000 initially					
	£1,000	.5	16%	£500	1.0	84%
	£0	.5				
2	Given £2,000 initially					
	-£1,000	.5	69%	-£500	1.0	31%
	£0	.5				

of and manipulating frequencies (e.g., Gigerenzer & Hoffrage, 1995; Sedlmeier & Betsch, 2002).

In DbS these three cognitive tools are used to derive the subjective value of an attribute value (i.e., to derive the psychological significance of a probability or of an outcome). Specifically, the subjective value is constructed from a series of binary, ordinal comparisons within a set of attribute values sampled both from the immediate context of the decision and from long-term memory. For each attribute value, a frequency count is kept of the number of favourable comparisons. The subjective value of an attribute is given by the proportion of favourable comparisons. For example, consider how the subjective value of a gain of £12 might be arrived at. The gain is compared to a small sample in working memory, say £1, £5, £33, £45, and £82. The sample will come from both the other attribute values in the immediate context and from previous experiences stored in long-term memory. A series of binary, ordinal comparisons are made between the target attribute of £12 and the attributes in the sample. Of the five possible comparisons, two of them are favourable (comparisons to £1 and £5). Thus the probability of a favourable comparison is 2/5, and in DbS this probability is used as the subjective value.

Of course, the subject value is completely dependent upon the distribution of attribute values in memory. Stewart et al. (2006) made the assumption that the distribution of attribute values in long-term memory reflects the real-world distribution (cf. J. R. Anderson & Schooler, 1991). Stewart et al. showed how the value and probability-weighting functions seen in prospect theory can be derived from the distributions of probabilities and amounts encountered in the environment, and thus provided an independent motivation for the particular shapes of these psychoeconomic functions.

The top panel of Figure 3 shows the credits into current accounts held by a large UK bank. The distribution is used as a crude approximation to the distribution of gains people encounter in the world. The figure is aggregated across many current accounts and represents about 320,000 payments in total. There are two key properties of this distribution. First, the distribution is very roughly power-law (i.e., is roughly linear on the log-log plot). Thus, although one might question whether the distribution of credits into UK current accounts is truly representative of the distribution of gains in an individual person's long-term memory, there is good reason to expect the long-term memory distribution to have similar
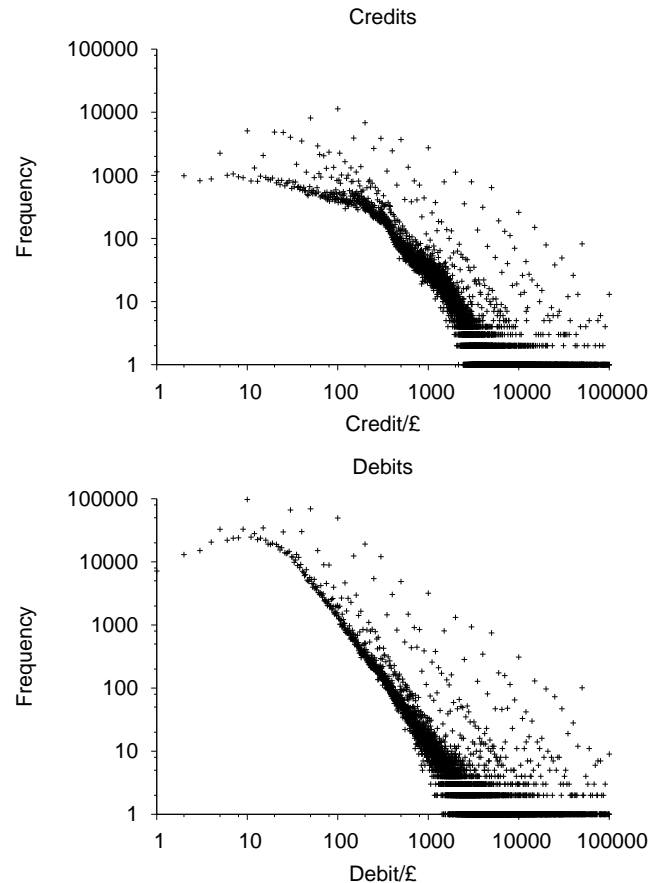

Figure 3. The frequency of credits to (top) and debits from (bottom) a large sample of UK current accounts. Adapted from Stewart, Chater, and Brown (2006).

properties. Power law distributions are very common (Bak, 1997) and describe the distributions of prices of many everyday items (e.g., Stewart et al., 2006; Stewart & Simpson, 2008). Second, and more importantly for the current argument, there are many more small credits than there are large credits. Thus, when one is sampling gains from long-term memory, one is more likely to sample small gains than large gains.

The subjective value of a given target gain is determined

by the proportion of gains in the sample to which a binary ordinal comparison is favourable. The proportion of a sample less than a given value is, of course, the definition of the cumulative probability density function. The top-right quadrant of Figure 4 shows the subjective value function for gains (or equivalently, the cumulative probability density function for gains). The subject value function for gains is concave, just like the utility function from EU theory and the value function for gains from prospect theory. But, in DbS, the concavity is derived from the interaction of a simple set of cognitive tools with the real-world distribution of attribute values. This contrasts strongly with EU theory and with prospect theory where the shape of the function is descriptive, that is, derived to fit the choice data from gambling experiments.

Previous work has estimated the exponent for power-law utility functions. Galanter (1962) has people estimate how much money would make them twice as happy as a reference amount. Assuming a power-law utility function, Galanter estimated the power to be 0.43. Kornbrot, Donnelly, and Galanter (1981) estimated the exponent using a signal detection procedure. By varying the (small) payoffs for hits, correct rejections, false alarms, and misses Kornbrot et al. (1981) estimate an exponent of 0.48. Galanter (1990) repeated his earlier procedure and found an exponent of 0.54. Galanter (1990) also reports an unpublished magnitude estimate experiment by Kornbrot which found an exponent of 0.43. Using the data in Figure 4, I estimate the exponent to be 0.47, which agrees well with the earlier figures. [2]

A given individual probably has only a sub-sample of the set of amounts in Figure 3 is their memory and, further, probably only samples a sub-sample of the amounts in their memory. Stewart et al. (2006) show that, under conditions of random sampling, Figure 4 actually represents the mean subject value. A more realistic assumption about sampling is likely to be necessary—sampling is surely not random—but, for the arguments made so far, more detailed assumptions are not necessary.

The bottom panel of Figure 3 and the bottom-left quadrant of Figure 4 repeat the credits analysis for debits. The distribution of debits from UK current accounts is also power law: There are more small debits than large debits. I have plotted the subjective value function for debits in the bottom-left quadrant of Figure 4. The function is reversed because large debits are worse than small debits (whereas large gains are better than small gains). Crucially, because there are more small debits than large debits, the subjective value function for is convex for losses, just as it is in prospect theory. Comparing credits and debits, there are more small debits than small credits, so the value function for losses is initially steeper than the value function for gains—just as in prospect theory—and offers an account of loss aversion. In summary, the DbS subjective value function is very similar to prospect theory's value function (see Figure 2). But in DbS the shape of the value function is derived from the interaction of a simple set of cognitive tools with the real-world distribution of gains and losses.

Stewart et al. (2006) repeated the credits and debits analysis for probabilities. People have a strong preference for
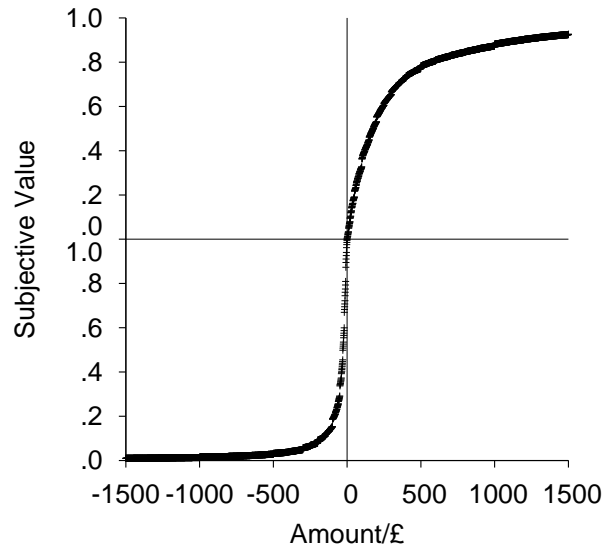

Figure 4. The DbS subjective value function derived from the distribution of credits and debits in Figure 3. Adapted from Stewart, Chater, and Brown (2006).

using words rather than numbers to express subjective probabilities (see Budescu & Wallsten, 1995, for a review). Figure 5 shows the frequencies in natural language with which different words are used to describe probabilities. Each line represents a particular word or phrase (e.g., "a fair chance"). The line's location on the abscissa gives the mean numerical probability the phrase is associated with, and the line's height represents the frequency in the British National Corpus.

The derivation of the subjective probability function is just the same as the derivation of the value functions. Figure 6 shows the subjective probability function (i.e., the proportion of verbal phrases describing events less likely than the target phrase and, equivalently, the cumulative probability density function). Because very unlikely probabilities (e.g., "impossible" and "never") and very likely probabilities (e.g., "always" and "definitely") are much more frequent than intermediate probabilities (e.g., "possible" and "fair chance"), the derived subjective probability function is most sensitive (i.e., steepest) at 0 and 1 and least sensitive (i.e., flattest) with intermediate probabilities. This DbS subjective proba-

[2] If data outside the range in Figure 4 are used, so that the DbS subjective value function is extended to cover the full range of values in Figure 3, the best-fitting exponent drops to 0.11: Plotted in log-log space, the subjective value function is initially linear with slope 0.47, but then the slope reduces to zero above about £10,000, so the subjective values of large amounts are all the same. This breakdown, I think, reflects the limits of using these current account data and assuming people randomly sample from them. For example, when considering an annual salary or the price of a house people are more likely to sample other similarly large amounts rather than the costs of cups of coffee or weekly shops. Scale invariance in the world and in memory (Chater & Brown, 2008) is likely to lead to similar shaped utility functions across a range of magnitudes.
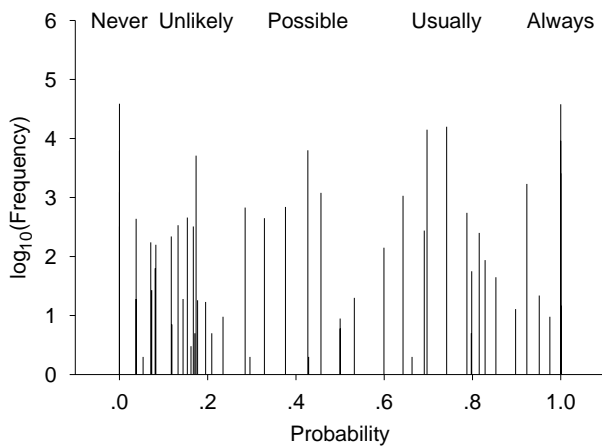
Figure 5. The frequencies of different probability phrases in the British National Corpus. Adapted from Stewart, Chater, and Brown (2006).
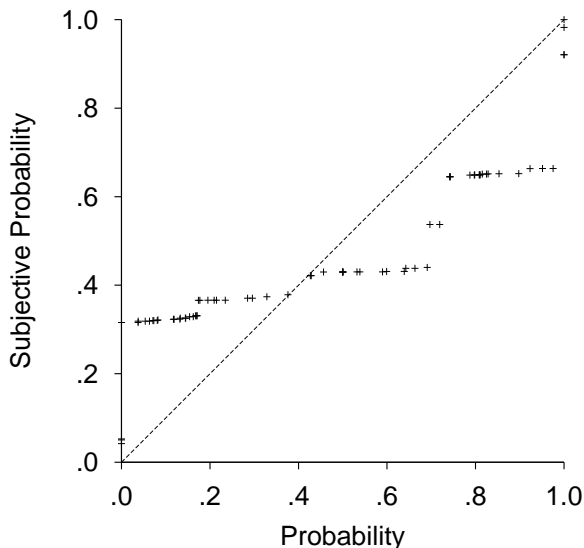

Figure 6. The DbS probability-weighting function derived from the distribution of probability phrases in the British National Corpus in Figure 5. Adapted from Stewart, Chater, and Brown (2006).

bility function is very similar to cumulative prospect theory's weighting function (see Figure 2). But in DbS the shape of the subjective probability function is derived from the interaction of a simple set of cognitive tools with the real-world distribution probabilities, just as it was for gains and losses.

DbS might well be extended to situations of uncertainty where the actual probabilities are not known (though I have not yet formulated the details of this extension or tested it). Because the construction of the subjective probability (or weight) of an event relies only on binary, ordinal comparisons between events, weights might be constructed even if one knows only which event is the more likely for pairs of events—knowing the actual probabilities is not required.

In summary, in DbS, the subjective value of a given attribute emerges from a series of binary, ordinal comparisons with a sample from long-term memory. The subjective value is given by the proportion of attribute values in the sample that are less favourable. As a result, the psychoeconomic functions are derived from the distribution of attribute values in the real world. These functions are in close agreement with the descriptive functions from prospect theory that describe people's decisions so well. Stewart et al. (2006) also apply DbS to temporal durations, and show how hyperbolic-like temporal discounting emerges from the distribution of delays that people encounter. More generally, in DbS all attribute values are treated in exactly the same way. For example, the DbS argument could be applied to give the subject value of any type of attribute (e.g., ipod capacities, broadband speeds, calorific values, etc.). This is a pretty serious departure from the normative consensus, where probabilities are treated differently from amounts (specifically, probabilities are used to weight amounts) but, as I'll describe below, the model provides a good description of the choices people make despite this departure.

A Model of Choices

Thus far, I have described the DbS account of the valuation of economic attributes: I've reviewed how the distribution of attribute values in the real world and the use of limited cognitive tools provides an independent motivation for the psychoeconomic functions that describe the choices we make so well. But of more interest is a model that actually makes choice predictions. What is needed is a mechanism for integrating information about risk and reward. For the models derived from EU theory, this integration is described by a multiplication. In DbS, mechanism for integration is additive.

Though there is some evidence for multiplicative integration in providing valuations of single risky prospects, there has been no comparison of different models of integration in choice. When people are asked to provide certainty equivalents, buying prices, and selling prices for gambles, their ratings tend to show an interaction between probability and amount information, indicating multiplicative information integration (e.g., Mellers, Chang, Birnbaum, & Ordóñez, 1992; Tversky, 1969). But ratings of the attractiveness of prospects tend to be additive (e.g., Levin, Johnson, Russo, & Delden, 1985; Mellers & Chang, 1994; Mellers, Ordóñez, & Birnbaum, 1992), integration of sample probabilities tends to be additive (e.g., Shanteau, 1975), and integration of attribute values tends to be additive for non-risky options (e.g., N. H. Anderson, 1981). To the best of my knowledge, no one has directly compared the fits of additive and multiplicative models of decision under risk to actual choices between pairs of risky options rather than to valuation or rating of single options. In sum, there does not seem to be strong empirical evidence for preferring multiplicative integration to additive integration in risky choices, though this is a controversial statement.

Stewart and Simpson (2008) have extended the DbS

model with a process account of the integration of information. The extension is quite simple: People are assumed to make a series of binary, ordinal comparisons between attribute values in working memory. Frequency accumulators tally the number of favourable comparisons for each option. A choice is made when the difference in tallies exceeds a threshold. For example, consider Choice 2 from Table 1. The choice offers (a) a 20% chance of £4000 otherwise an 80% chance of £0 or (b) a 25% chance of £3,000 otherwise a 75% chance of £0. The decision process is hypothesised to proceed as follows:

1. A target attribute is randomly selected by selecting a gamble, an attribute type, and an attribute at random. For example, Prospect A might be selected (rather than Prospect B), amounts might be selected (rather than probabilities), the value £3000 might be selected (rather than £0).

2. A comparison attribute is randomly selected from the decision sample. The decision sample will comprise both attribute values from the immediate context (i.e., those listed for Choice 2) and attribute values from long-term memory (approximated by the distributions from Figure 3). Stewart and Simpson (2008) assumed sampling from either source is equally likely. Let's say a value of £2000 (from long-term memory) is sampled.

3. The target and comparison attribute values are compared with a binary, ordinal comparison. In the current example, the comparison is between a target value of £3000 and a comparison value of £2000. £3000 is good compared to £2000, and thus the comparison is favourable.

4. If the comparison is favourable, the accumulator for the target prospect is incremented by one count. In this way, the accumulators tally the number of favourable binary, ordinal comparisons.

5. If the difference between accumulator tallies for each prospect reaches threshold, select the prospect with the highest accumulator. Otherwise, begin again at Step 1.

For a given option, a favourable comparison involving one of its probabilities and a favourable comparison involving one of its amounts both lead to an increment for the option's accumulator. Effectively then, information about risk and information about reward are combined additively. This additive combination is a break from the multiplicative combination in the EU-theory based models, but nonetheless provides a good account of people's risky choices.

The probability that the accumulator of a given option will be incremented is related in a very straightforward way to the original DbS model subjective values. For a given attribute value, the probability of a favourable comparison is given by the proportion of attribute values in working memory that are less favourable. If one assumes that any attribute value of an option is equally likely to be selected for comparison, then the probability of an increment of the frequency accumulator for that option is given by the average, across all of the option's attribute values, of the proportion of less favourable attribute values in working memory. In other words, the probability that an option's accumulator is incremented is given by the average subjective value of the option's attributes.

A note about comparing probabilities is in order. Con-

sider the comparison between 80% in Prospect B and 25% in Prospect A. Just because 80% is greater than 25% does not mean the comparison is favourable. After all, an 80% chance of nothing is clearly worse than a 25% chance of something. Effectively, the valence of the corresponding outcome needs to be considered when comparing probabilities. In the mathematics of the model, if the corresponding outcome is bad, the accompanying probability is given a negative sign, so that probabilities of bad things, no matter how large, are always less favourable than the probabilities of good things, no matter how small. (This point will be important later on when considering why DbS works.)

One strength of this process version of the DbS model is that a closed form mathematical expression for the probability of selecting either option can be formed. Because, in the case of binary choice, one is incrementing either the accumulator for one option or the accumulator for the other and because the choice rule is a difference threshold, the model can be implemented using the mathematics of the random walk (Feller, 1968). Thus there is no need to run simulations of the model. It is sufficient to calculate the average proportion of less favourable attribute values for a given option, and use these as drift rates in the random walk. Specifically, the probability of choosing Prospect A is $1/(1 + (1 - 1/d)^T)$ where d is the relative probability of an increment for Prospect A compared to Prospect B and T is the threshold. Alternatively the drift rates can drive a race to a fixed, absolute threshold, in which case the probability of choosing Prospect A is $\sum_{i=0}^{T-1} \frac{(T-i-1)!}{i!(T-1)!} p_A{}^T p_B{}^i$ where p_A is the probability of an increment for Prospect A and $p_B = 1 - p_A$ is the probability of an increment for Prospect B. Stewart and Simpson (2008) found that both implementations work well.

Stewart and Simpson (2008) have shown how the model provides a good account of the choice proportions from the original Kahneman and Tversky (1979) choices, including the common-ratio and framing effects described earlier. Figure 7 reproduces Stewart and Simpson's plot of the model's predictions of the choice proportions for each option. Each point in the plot represents a choice. The y-axis gives the empirical choice proportion for selecting (arbitrarily) the right-hand option that Kahneman and Tversky found, and the x-axis gives the DbS model prediction for the probability of selecting the right option. There is good agreement between the model predictions and the data ($r^2 = .87$), and this is mainly due to the model predicting the correct direction of preference for each prospect (i.e., the points fall in either the top-right or bottom-left quadrants).

In fitting this data, the DbS model did not have any free parameters. Stewart and Simpson (2008) explore how robust these predictions are under alternative implementations of the model (with a free parameter representing the probability of selecting amounts rather than probabilities for sampling, with a free parameter representing the relative weighting of the background distribution of attribute values and the attribute values from the immediate context in the decision sample, with uniform rather than skewed distributions of attribute values, and with alternative thresholds or alter-
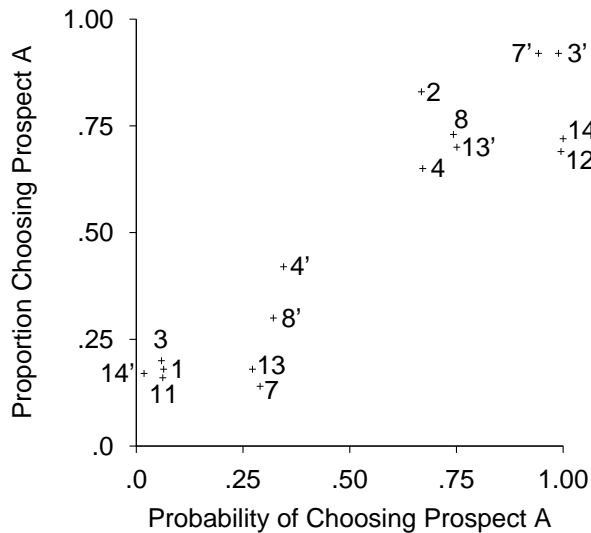
Figure 7. Choice proportions from the Kahneman and Tversky (1979) data set plotted against DbS predictions. Data point numbers match Kahneman and Tversky's numbering. Adapted from Stewart and Simpson (2008).

native stopping rules). For a wide range of alternative parameter values, the model makes the same qualitative predictions, correctly predicting the direction of preference for all 16 choices. This is important because it suggests that the model is not performing well because it is too flexible: Stewart and Simpson (2008) were not just lucky in their assumptions about sampling, weightings of background and immediate context, etc.

There is, in fact, a good a priori reason to expect the DbS model to perform well. EU theory and its derivative models are, effectively, regression models. For example, consider a n-outcome gamble of the form p_1 chance of x_1, p_2 chance of x_2, ..., p_n chance of x_n. According to these models, the subjective value of the gamble is given by $w(p_1)U(x_1) + w(p_2)U(x_2) + ... + w(p_n)U(x_n)$. The w and U functions transform probabilities and amounts, respectively, into their subjective equivalents. The subjective probabilities are effectively used as regression coefficients to weight the utilities in arriving at an overall subjective value for the gamble. Equivalently, because multiplication is commutative, the utilities can be thought of as regression coefficients weighting the subjective probabilities. (This unusual interpretation comes into play below.) Dawes (1979) has shown how improper linear models provide a very good approximation to linear regression equations. In an improper linear model, the magnitude of the regression coefficients is dropped and only their sign is retained. For example, for the regression equation $3x_1 - 4x_2 + 7x_3 - 2x_4$, the improper counterpart is $+x_1 - x_2 + x_3 - x_4$. Though the cognitive process in the DbS model is not an improper linear regression equation—the process is a random walk—the predictions of the model are quite similar to those of an improper linear model. Because

of the way that DbS compares the probabilities of good and bad things, probabilities are weighted with either $+1$ or -1 depending on the valence of the associated outcome. Thus, as a regression equation, the subjective value of a gamble (i.e., the probability of an increment to the associated accumulator) is given by $val(x_1)R_p(p_1) + val(x_2)R_p(p_2) + ... + val(x_n)R_p(p_n) + R_x(x_1) + R_x(x_2) + ... + R_x(x_n)$ where $val(x_i)$ is either $+1$ or -1 depending on the valence of x_i, and R_p and R_x are functions giving the proportion of attribute values in the sample that are less favourable than the target attribute value. The $val(x_i)R_p(p_i)$ terms in the DbS model are the improper version of the $w(p_i)U(x_i)$ terms in the EU-based models. Thus just as utilities act as regression coefficients for subjective probabilities in the EU-based models, so the valence of the amount acts as an improper regression coefficient for the subjective probability in the DbS model. So, because the EU-based models provide a good description of people's choices, to the extent that the improper approximation is good, the DbS model should also provide a good description of people's choices.

Experimental Evidence

Thus far, I've described how, in DbS, the distribution of attribute values in the environment combines with the use of a limited set of cognitive tools to offer an account of why the psychoeconomic functions inferred from our risky decisions take the forms that they do. I've also reviewed how DbS might be extended to predict risky choices and how, on a preliminary test, this extension seems able to account for the now infamous violations of EU theory that were used to motivate prospect theory. Of course, there are many more important results in risky decision making that a complete model must account for and this is work currently in progress. I close this article with a review of some of the experimental evidence that motivated the DbS model. In each of the following sections the experimental data provide a challenge to the notion that our subjective valuations of outcomes and probabilities are stable. Instead, the data are consistent with the DbS model, which predicts that subjective valuations will vary as the distribution of attribute values in the immediate context changes.

Prospect Relativity

Birnbaum (1992) and Stewart, Chater, Stott, and Reimers (2003) asked participants to select the certainty equivalent for a prospect from a series of candidate values. The certainty equivalent for a risky prospect is the amount of money available with certainty that is worth the same as the chance to play the risky prospect. For example, £40 would be the certainty equivalent for the gamble 50% chance of £100 if people were indifferent between receiving £40 or playing the 50% chance of £100 gamble. Birnbaum manipulated the skew of the candidate certainty equivalents. When the candidate values were positively skewed (i.e., many small values) the certainty equivalent selected was smaller than when the candidate values were negatively skewed (i.e., many large values). Stewart et al. manipulated the range of candidate

certainty equivalents. When the range was high (i.e., all certainty equivalents were large) the prospect was overvalued compared to when the range was low (i.e., all certainty equivalents were small). Instead, participants were selecting certainty equivalents in the middle of the range, irrespective of the absolute value of the certainty equivalents and despite that fact that failing to provide their true certainty equivalents was costing them money. In sum, these experiments show that the distribution of candidate certainty equivalents affects their subjective valuation (see also Ariely, Koszegi, Mazar, & Shampan'er, n.d.).

EU-based theories, in which a utility or value function (e.g., Figure 1) is used to transform amounts into their subjective equivalent, cannot account for this result: EU-based theories assume that the transformation is stable, but these data suggest that the utility of a candidate value varies depending on the accompanying candidate values. And this is just was DbS predicts: The subjective value of a certainty equivalent is determined by binary, ordinal comparison with the sample in working memory—which is likely to be full of the suggested certainty equivalents. Thus a given certainty equivalent will seem subjectively larger when there are many smaller candidate certainty equivalents in the choice set. But the same certainty equivalent will seem subjectively smaller when there are many larger candidate certainty equivalents in the choice set.

Stewart et al. (2003) conducted an accompanying choice experiment showing that prospects are also valued relative to one another. Participants were asked to select the prospect they'd most like to play from a set where risk and reward were traded off. In the High-Risk Condition one group of participants chose from five relatively high-risk prospects: {50% chance of £50, 55% chance of £45, 60% chance of £40, 65% chance of £35, 70% chance of £30}. In the Low-Risk Condition another group participants chose from a set of relatively low-risk prospects: {75% chance of £25, 80% chance of £20, 85% chance of £15, 90% chance of £10, 95% chance of £5}. The high-risk prospects can be derived from the low risk prospects by decreasing all of the probabilities by a fixed value and increasing all of the amounts by a fixed value. Because DbS assumes that attributes are valued by comparing them to other attribute values in the choice set, decreasing all of the probabilities by a fixed value or increasing all of the amounts by a fixed value will not affect the subjective values of the prospects. Thus DbS predicts a similar pattern of preference in the two conditions, which is just what Stewart et al. (2003) found. In contrast, according to EU theory, the even distribution of preferences across the Low-Risk Condition means that people generally have quite a low risk preference and thus the lowest-risk option in the High-Risk Condition should be really popular. Similarly, the even distribution of preferences across the High-Risk Condition means that people have quite a high risk preference and thus the highest-risk option in the Low-Risk Condition should be really popular. Thus EU theory cannot predict an even pattern of preference in both conditions. These data suggest that the subjective value of a prospect is derived relative to the set of accompanying prospects. Benartzi and Thaler (2001) and

Vlaev, Chater, and Stewart (2007a, 2007b) give a series of real-world examples of prospect relativity.

The Attraction Effect

The attraction effect offers another example of how preference between risky prospects can be altered as the choice set is manipulated. The attraction effect is extremely well replicated across a wide variety of stimulus attributes (e.g., Huber, Payne, & Puto, 1982; Dhar & Glazer, 1996; Simonson, 1989), but here I concentrate on decision under risk where Wedell (1991) has explored the effect in detail. Wedell (1991) offered participants choices between a low-probability-high-amount prospect and a high-probability-low-amount prospect. In the example in Figure 8, participants might choose between Prospect A (an 83% chance of $12) and Prospect B (a 30% chance of $33). The third gamble in the set is manipulated between choices. In each case, the added prospect is clearly worse—offering a slightly lower amount with a slightly lower probability—than one of the original two prospects. In Figure 8 Prospect C (a 78% chance of $10) is dominated by Prospect A, and Prospect D (a 25% chance of $30) is dominated by Prospect B. In each case, participants have a preference for the dominating option: When the choice set is {Prospect A, Prospect B, and Prospect C} participants prefer Prospect A. When the choice set is {Prospect A, Prospect B, and Prospect D} participants prefer Prospect B. In sum, the manipulation of a dominated option causes a preference reversal between Prospects A and B. This results in a challenge to the classic account, in which Prospects A and B are valued independently of one another and independently of any other prospects (i.e., Prospects C and D).

A DbS account of the basic attraction effect is quite straightforward. Comparisons between the dominated option and the dominating option favour the dominating option for both probability and amount. But comparisons between the dominated option and the non-dominating option favour the dominated option on one attribute and the non-dominating option on the other attribute. Thus the probability of a favourable comparison is raised more for the dominating option. The DbS account is essentially the same as the account offered by other models in which the effect results from the dominated alternative altering the subjective values of the other two options (e.g., range-frequency theory, described later).

Wedell (1991) provided further detailed results by manipulating the location of the dominating option (see also Huber et al., 1982; Dhar & Glazer, 1996). Some of this evidence appears problematic for value-shifting accounts including the DbS account above, so conclusions about the adequacy of the DbS explanation must await further work.

Probability Judgements, Working Memory, and Context Effects in Probability Judgement

DbS predicts a close link between judgement and working memory. Working memory capacity reflects the quantity of information that can be held in mind whilst com-
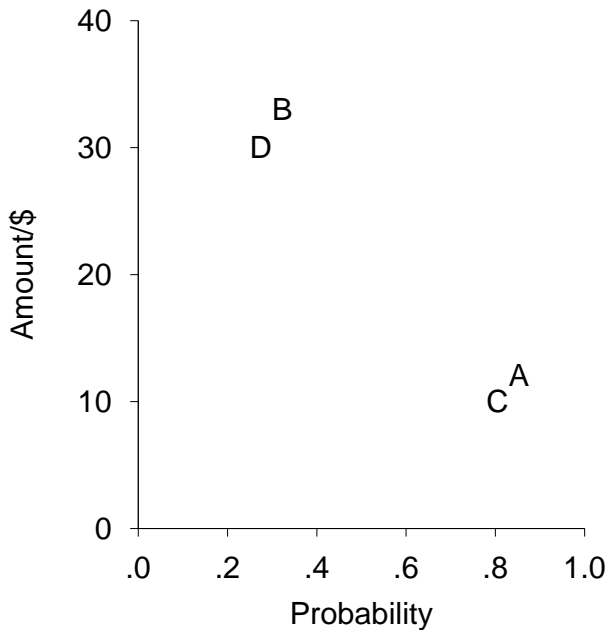
Figure 8. Some options used by Wedell (1991) to demonstrate the attraction effect.

pleting other cognitive processes (Engle, Tuholski, Laughlin, & Conway, 1999, cf. passive short-term memory capacity). Dougherty and Hunter (2003a, 2003b) and Sprenger and Dougherty (2006) find that probability judgements are overestimated more by individuals with lower working-memory capacities. The explanation is that lower working memory capacity means that fewer alternative hypotheses can considered, and thus the target hypothesis is judged more likely. Thomas, Dougherty, Sprenger, and Harbison (2008) give a review and a mathematical model; here, this finding is taken as evidence that event probabilities are compared to a sample of alternative event probabilities taken from memory.

There is also evidence that the distribution of the event probabilities in working memory affects these judgements. Windschitl and Chambers (2004) found that adding an extremely unlikely alternative to the set of events increases judgements of the probability of the target event. For example, when told that one of the cities Calcutta, Cincinnati, Nairobi, or Moscow lies below the equator and asked to judge subsequently the probability that Nairobi is below the equator, including Cincinnati and Moscow (which are extremely unlikely to lie below the equator) increases judgements. Windschitl and Chambers explain these data using a contrast mechanism, in which the probability for Nairobi is inflated because the probabilities for Moscow and Cincinnati are so small.

Even more direct evidence comes from Windschitl and Wells (1998). Windschitl and Wells asked participants to judge the likelihood that they would win a lottery if they held a given number of tickets. Subjective judgements were affected by the distribution of tickets among other players.

For example, holding 21 tickets was judged more favourably when other players held 15, 14, 13, 13, and 12 tickets than when other players held 52, 6, 2, 2, and 5 tickets. (Control conditions address the possibility that judgements differed because people failed to sum the total number of other tickets correctly.) Windschitl and Wells attribute these effects to a contrast mechanism, and DbS is just this type of mechanism. In DbS, comparing 21 to every item in the first set results in five favourable comparisons. But comparing 21 to every item in the second set results in only four favourable comparisons. Thus the subjective probability associated with 21 is higher in the first set. In summary, for probability judgement, there is good evidence for the comparisons with a sample of event probabilities in working memory.

Salary Satisfaction

How happy you are with your salary does not just depend on how much you earn. It depends on how much you earn compared to your peers. Brown, Gardner, Oswald, and Qian (2008) examined the relationship between reported salary satisfaction and the distributions of salaries experienced. Brown et al. used a large data set completed by a sample of employees from a sample of UK companies. Because many employees in each of many companies were sampled, it was possible to investigate the effects of the distribution of salaries within a given employee's company had on the employee's satisfaction. In a regression analysis, after removing the effects of absolute salary, Brown et al. found a significant effect of (a) the rank of the employee's salary within the company and (b) the position of the employee's salary relative to the minimum and maximum salaries in the company. In summary, judgements of satisfaction with salary are higher if it is among the highest in the company, independently on the absolute level of the salary. And, in further analysis, Brown et al. found that quit rates were higher in companies where salary distributions were more positively skewed (i.e., many small salaries, few large salaries). Thus not only ratings of satisfaction but actual decisions to quit one's job depend on the distribution of salaries: A given salary does not map to a specific utility; instead, the utility of a given salary depends on how it compares to the other salaries one thinks about.

Revealed Psychoeconomic Functions

As I've reviewed above, the utility and subjective probability functions from EU theory and prospect theory are descriptive: They take the forms they do because they describe the risky choices we make. Typically, we infer the shapes of these functions by fitting models with free parameters to choice data (e.g., Gonzalez & Wu, 1999). It is even possible to discriminate between different candidate functional forms for utility and subjective probability functions using this technique (Stott, 2006).

But, according to DbS, the functions that are revealed will depend upon the distribution of attribute values that people encounter. Although it is not possible to manipulate experimentally the background distributions that people experience in their everyday lives, it is possible to manipulate the

distribution of attribute values that people experience in the laboratory when they are making the choices that will be used to infer their psychoeconomic functions. Stewart and Reimers (2008b) have done just this. In their experiment, participants made a series of choices between prospects. The prospects were created by crossing the probabilities .2, .4, .6., .8 and 1.0 with a set of amounts. The set of amounts was manipulated between participants and was either positively skewed {£10, £20, £50, £100, £200, £500} or negatively skewed {£10, £310, £410, £460, £490, £500}. The negatively skewed set is the mirror image of the positively skewed set.

To reveal the utility functions, Stewart and Reimers (2008b) fitted EU theory to the data. Rather than fitting a specific functional form for the utility function, they left the utility of each monetary amount as a free parameter. Figure 9 shows the best fitting (non-parametric) utility functions for the positively skewed and negatively skewed conditions. The most striking result is that the manipulation of the distribution of amounts people experienced produces quite different utility functions. The dashed lines illustrate the DbS predictions if the utility of each monetary amount was given by the proportion of smaller amounts in the set of amounts used in the experiment. For the positive skewed condition, the data match these predictions quite closely. For the negatively skewed condition, the utility function is more linear than the prediction. This actually makes sense if one assumes that the sample of amounts in working memory is a mixture of attributes from long-term memory and attributes from the immediate context as DbS assumes. The distribution in long-term memory is positively skewed (as expected if the distribution in memory represents the real-world distribution, Stewart et al., 2006). So, in the positively skewed condition, the overall decision sample should be positively skewed because both the long-term-memory and the immediate-context distributions are positively skewed. But in the negatively skewed condition, the overall decision sample should be a mixture of the positively-skewed long-term-memory distribution and the negatively-skewed immediate-context distribution—and together this mixture should be closer to uniformly distributed which would in turn predict a more linear utility function.

In summary, experimentally manipulating the distribution of attribute values produces differences in the revealed utility functions. This finding is problematic for the notion that utility functions are stable psychological entities. In the traditional view, in a given situation, the utility of each outcome is looked up and then the option with the higher EU is chosen. To accommodate Stewart and Reimers's (2008b) data within this framework, the best scenario is one in which the functions are assumed to vary from context to context. But then a theory of how the functions vary with experimental context in needed to provide explanatory power. However, these results follow naturally from DbS. But the distribution of attribute values in working memory will vary over time—because of both stochasticity in sampling and differences in the immediate distribution of attribute values—the subjective value of a given attribute will vary from context to context
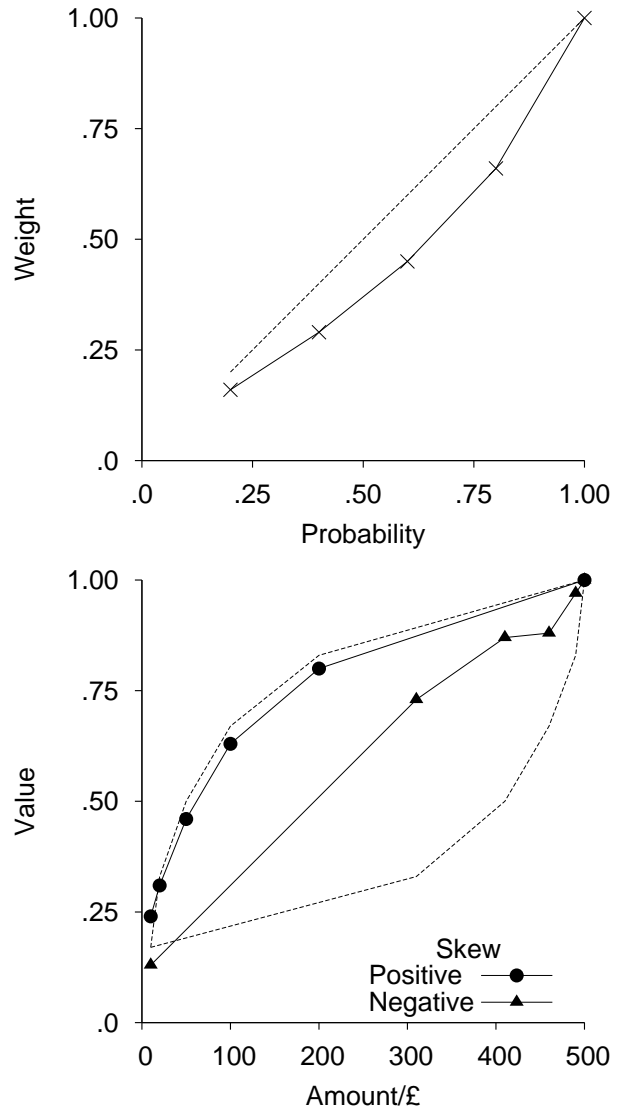

Figure 9. Revealed utility functions adapted from Stewart and Reimers (2008b). Solid lines represent data and dashed lines represent DbS predictions (with only immediate attribute values in the decision sample).

and time to time.

Preference Reversals Induced by Attribute Distribution Differences

As I described when I reviewed the common-ratio and framing effects above, preference reversals are particularly important in the decision-making literature. Stewart and Reimers (2008a) manipulated the distribution of attribute values in the immediate context to create a preference reversal in a critical choice. The critical choice was between a 30% chance of 100 ipoints and a 40% chance of 75

ipoints.[3] To set the context, this critical choice was preceded by eight choices in which the distribution of attribute values was manipulated between participants. In the Probabilities-Together-Amounts-Apart Condition, the probabilities 10%, 20%, 30%, 40%, 50%, and 60% were crossed with amounts 75 ipoints, 80 ipoints, 85 ipoints, 90 ipoints, 95 ipoints, and 100 ipoints to create a set of prospects (100 ipoints are worth about £1). In the Probabilities-Apart-Amounts-Together Condition, the probabilities 30%, 32% , 34%, 36%, 38%, and 40% were crossed with the amounts 25 ipoints, 50 ipoints, 75 ipoints, 100 ipoints, 125 ipoints, and 150 ipoints.

The distribution of attributes in the Probabilities-Together-Amounts-Apart Condition was selected to make the difference between the 30% and 40% in the critical choice appear small (with probabilities ranking 3rd and 4th in the distribution) while the difference between the 75 ipoints and 100 ipoints was made to appear large (with amounts ranking 1st and 6th in the distribution). Thus, because the difference in amounts is subjectively much larger than the difference in probabilities, people should select on the basis of amount and choose the 30% chance of 100 ipoints. The Probabilities-Apart-Amounts-Together Condition reversed this manipulation, with probabilities ranking 1st and 6th whilst amounts ranked 3rd and 4th. In this condition, people should select on the basis of probability and select the 40% chance of 75 ipoints.

Figure 10 shows how the proportion of people selecting each option in the critical choice varied by condition in just this way. The switch from a majority preference for a 30% chance of 100 ipoints in the Probabilities-Together-Amounts-Apart condition to a majority preference for a 40% chance of 75 ipoints in the Probabilities-Apart-Amounts-Together Condition represents an attribute-distribution induced preference reversal. Again, the conclusion is that we do not have stable underlying psychoeconomic functions. The best case interpretation for the classical view is that these functions are malleable and vary from context to context, but to go beyond the merely descriptive, one needs a theory to explain why these functions vary from context to context. DbS provides an account of this sort - by abandoning stable psychoeconomic functions and instead assuming that subjective values are constructed afresh for each preference using simple cognitive tools.

Discussion

In the review, I have presented evidence that the subjective value of a given risky option is not derived independently of the other options on offer. DbS offers one account of why this might be the case. Because the decision sample contains attributes from the immediate context, the subjective value of each attribute value will vary from context to context as the contents of the decision sample varies from context to context. In this way, DbS makes reference to three significant bodies of work: Parducci's range-frequency theory, Poulton's response-contraction explanation of prospect theory, and Payne, Bettman, Johnson, Slovic and Luce's construction-of-preference concept.
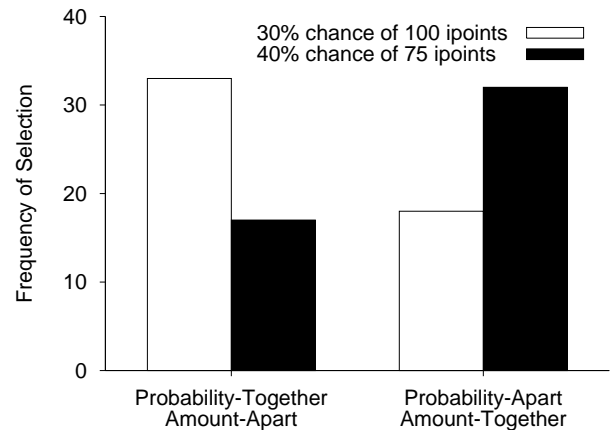

Figure 10. A preference reversal for the choice 30% chance of 100 or 40% chance of 75 ipoints Stewart and Reimers (2008b).

Range-Frequency Theory. Range-frequency theory (Parducci, 1965, 1995) is a theory of category judgement. In category judgement tasks, stimuli varying along a single dimension are assigned one label from an ordered set of categories (e.g., lines varying in their length might be assigned to categories "very short", "short", "medium", "long", "very long"). Originally, the theory was used to account for the effect of the distribution of perceptual stimuli on category judgements. For example, a given target line length is rated as larger in a positively skewed distribution of line lengths (i.e., many smaller lines) than in a negatively skewed distribution of line lengths (i.e., many larger lines). Range-frequency theory has been applied very widely beyond the categorisation of perceptual stimuli (see Parducci, 1995, for a review).[4]

Range-frequency theory has two components. The range principle states that the stimulus range is divided into equal size categories (one for each category label) irrespective of the distribution of stimuli. The frequency principle states that the stimulus range is divided into categories so that each category is used equally frequently and contains an equal number of stimuli. Thus the division of the stimulus range under the frequency principle is completely dependent on the distribution of stimuli. For example, if line lengths are positively skewed, then the smaller lengths will be divided into many categories and the larger lengths into fewer categories. Effectively, under the frequency principle, the category label assigned to a particular stimulus is determined by the stimulus's rank position. The overall category assigned to a given

---

[3] ipoints is an online reward scheme. ipoints can be redeemed for a large range of goods.

[4] Adaptation level theory (Helson, 1964) represents perhaps the first attempt to account for contextual effects. In adaptation level theory, stimuli are judged against the mean of the distribution in which they are encountered. Range-frequency theory goes further in accounting for the effects of higher moments, like the variance and the skew.

stimulus is a weighted average of the categories given by the range and frequency principles.

DbS makes very similar predictions, but the theoretical account is quite different, offering a process rather than a description. Recall that the subjective value of an attribute is constructed from a series of binary, ordinal comparisons with attributes in the decision sample. The long-term memory contribution to the decision sample provides a component of the subjective value that is independent of the distribution of attribute values in the immediate context (cf. the range principle). The immediate context's contribution to the decision sample provides a component of subjective value that is dependent on the rank position of the target attribute value in the immediate context distribution (cf. the frequency principle). So although theoretical motivations are very different— equal division of the stimulus range and equal use of categories in range-frequency theory and construction of a subjective value via a series of binary, ordinal comparisons in DbS—the net effect is very similar. DbS goes beyond range-frequency theory in predicting that psychoeconomic functions emerge from the real-world distribution of attribute values and in providing mechanism for integrating information across attributes.

Response Contraction

Poulton (1994) offers an account of prospect theory in terms of response contraction. Response contraction is a bias of quantifying judgements in which responses are assimilated towards a reference point on the scale, often the midpoint of the scale (Poulton, 1989). Such a contraction provides a natural account of Kahneman and Tversky's probability weighting function in which small probabilities are overweighted as they are assimilated towards the centre of the scale and large probabilities are underweighted as they too are assimilated towards the centre of the scale. Similarly for gains and losses, although response contraction does not offer an account of the curvature of the value function, it does predict that large gains and large losses will be underweighted as they are assimilated towards the centre of the scale.

Poulton's (1994) response contraction explanation of these effects differs from the DbS account. Under DbS, the revealed psychoeconomic functions take the forms they do because of the distribution of attribute values people encounter. But at a more general level there is a common thread. DbS was derived from my work on contextual effects perceptual identification and categorisation (e.g., Stewart, Brown, & Chater, 2002; Stewart & Brown, 2004; Stewart et al., 2005; Stewart & Morin, 2007) and from classic studies in psychophysics (e.g., Garner, 1954). The hypothesis is that our representations of economic quantities like money, risk, and delay are similar in quality to our representation of perceptual quantities like loudness or brightness. In this respect, the argument follows Poulton's general argument that biases in judgement and decision making have the same causes as biases in other judgements. I have tried to go further in offering ideas about the common cognitive mechanisms responsi-

ble for the common biases.

The Construction of Preference

In DbS, the psychoeconomic functions do not actually have a psychological reality. That is, the functions are not internal to the individual, and carried about with them ready to apply to each new decision they make. So under DbS the psychoeconomic functions can be considered as revealed from or descriptive of the choices that people make given the real-world distribution of attribute values. But these functions do not describe the psychological processes that underlie them. That is, people do not simply use these functions to "look up" the psychological value of a given gain, loss, or probability. Instead, the psychological process is a series of binary, ordinal comparisons with a sample of attribute values and the subjective value is the endpoint of an accumulation of favourable comparisons. Further, the subjective value will vary as a function of the real-world distribution of attribute values and as a result of stochasticity in sampling of these values. In this respect, DbS makes a connection with the literature on the construction of preference (cf. Bettman, Luce, & Payne, 1998; Payne, Bettman, & Johnson, 1992; Slovic, 1995)—the notion that our preferences are not stable but are derived afresh for each new decision.

Conclusion

I have argued that the decisions that we make are the result of our use of a set of domain-general cognitive tools. In the absence of stable functions relating risks and rewards to their subjective equivalents, we value risks or rewards by comparing them with samples of risks and rewards from memory. And because our memories are adapted to represent the environment, the memory samples reflect the real-world distributions of attribute values that we encounter. Psychoeconomic functions describing our valuation of risks and rewards can be derived from the distribution of attribute values we encounter in the real world, though these functions do not themselves describe the psychological processes underlying their derivation.

Risks and rewards are valued by the same binary, ordinal comparison mechanism because, without formal training, we do not have any other method for assessing the importance of these numbers. So, although probabilities and amounts should be treated differently in a normative account, they are processed in the same way as one another (and other psychoeconomic attributes) in this psychological account.

A key test of the DbS hypothesis is that experimental manipulation of the distributions of attribute values should have a strong effect on the decisions that we make. I have presented evidence that manipulating the distribution does indeed have this effect. Though some have argued that these effects should be explained away, averaged over, or carefully counterbalanced out, I think they indicate the true base for decision under risk. My conclusion is that decision is not by "look up"; decision is by sampling.

References

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.

Anderson, J. R., & Schooler, L. J. (1991). Reflections of the environment in memory. *Psychological Science*, *2*, 396–408.

Anderson, N. H. (1981). *Foundations of information integration theory*. New York: Academic Press.

Ariely, D., Koszegi, B., Mazar, N., & Shampan'er, K. (n.d.). *Price-sensitive preferences*. Unpublished manuscript.

Bak, P. (1997). *How nature works: The science of self-organized criticality*. Oxford, UK: Oxford University Press.

Bell, D. E., & Fishburn, P. C. (1999). Utility functions for wealth. *Journal of Risk and Uncertainty*, *20*, 5–44.

Benartzi, S., & Thaler, R. H. (2001). Naive diversification strategies in defined contribution saving plans. *The American Economic Review*, *91*, 79–98.

Bernoulli, D. (1954). Expositions of a new theory of the measurement of risk. *Econometrica*, *22*, 23–36. (Original work published 1738)

Bettman, J. R., Luce, M. F., & Payne, J. W. (1998). Constructive consumer choice processes. *Journal of Consumer Research*, *25*, 187–217.

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

Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review*, *115*, 453–501.

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

Brown, G. D. A., Gardner, J., Oswald, A. J., & Qian, J. (2008). Does wage rank affect employees' well-being? *Industrial Relations*, *47*, 355–389.

Budescu, D. V., & Wallsten, T. S. (1995). Processing linguistic probabilities: General principles and empirical evidence. In J. Busemeyer, D. L. Medin, & R. Hastie (Eds.), *The psychology of learning and motivation: Vol. 32. Decision making from a cognitive perspective* (pp. 275–318). San Diego, CA: Academic Press.

Busemeyer, J. R., & Townsend, J. T. (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 J. Kagel & A. E. Roth (Eds.), *Handbook of experimental economics* (pp. 587–703). Princeton, NJ: Princeton University Press.

Carlin, P. S. (1992). Violations of the reduction and independence axioms in Allais-type and common-ratio effect experiments. *Journal of Economic Behavior & Organization*, *19*, 213–235.

Chater, N., & Brown, G. D. A. (2008). From universal laws of cognition to specific cognitive models. *Cognitive Science*, *32*, 36–47.

Daniels, R. L., & Keller, L. R. (1992). Choice-based assessment of utility functions. *Organizational Behavior and Human Decision Processes*, *52*, 524–543.

Davies, G. B., & Satchell, S. E. (2007). The behavioural components of risk aversion. *Journal of Mathematical Psychology*, *51*, 1–13.

Dawes, R. M. (1979). The robust beauty of linear models of decision making. *American Psychologist*, *34*, 571–582.

Dhar, R., & Glazer, R. (1996). Similarity in context: Cognitive representation and violation of preference and perceptual invariance in consumer choice. *Organizational Behavior and Human Decision Processes*, *67*, 280–293.

Dougherty, M. P. R., & Hunter, J. (2003a). Hypothesis generation, probability judgment, and individual differences in working memory capacity. *Acta Psychologica*, *113*, 263–282.

Dougherty, M. P. R., & Hunter, J. (2003b). Probability judgment and subadditivity: The role of working memory capacity and constraining retrieval. *Memory & Cognition*, *31*, 968–982.

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

Engle, R. W., Tuholski, S. W., Laughlin, J. E., & Conway, A. R. A. (1999). Working memory, short term memory, and general fluid intelligence: A latent variable approach. *Journal of Experimental Psychology: General*, *128*, 309–331.

Ert, E., & Erev, I. (2007). *Loss aversion in decisions under risk and the value of a symmetric simplification of prospect theory*. Unpublished manuscript.

Feller, W. (1968). *An introduction to probability theory and its applications* (Vol. 1). New York: Wiley.

Fishburn, P. C., & Kochenberger, G. A. (1979). Two-piece von Neumann-Morgenstern utility functions. *Decision Sciences*, *10*, 503–518.

Galanter, E. (1962). The direct measurement of utility and subjective probability. *American Journal of Psychology*, *75*, 208–220.

Galanter, E. (1990). Utility functions for nonmonetary events. *American Journal of Psychology*, *103*, 449–470.

Garner, W. R. (1954). Context effects and the validity of loudness scales. *Journal of Experimental Psychology*, *48*, 218–224.

Gigerenzer, G., & Hoffrage, U. (1995). How to improve Bayesian reasoning without instruction: Frequency formats. *Psychological Review*, *102*, 684–704.

Gonzalez, R., & Wu, G. (1999). On the shape of the probability weighting function. *Cognitive Psychology*, *38*, 129–166.

Helson, H. (1964). *Adaptation-level theory*. New York: Harper & Row.

Huber, J., Payne, J. W., & Puto, C. (1982). Adding asymmetrically dominated alternatives: Violations of regularity and the similarity hypothesis. *Journal of Consumer Research*, *9*, 90–98.

Kahneman, D., & Miller, D. T. (1986). Norm theory: Comparing reality to its alternatives. *Psychological Review*, *93*, 136–153.

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

Kahneman, D., & Tversky, A. (Eds.). (2000). *Choices, values and frames*. New York: Cambridge University Press & the Russell Sage Foundation.

Kornbrot, D. E., Donnelly, M., & Galanter, E. (1981). Estimates of utility function parameters from signal-detection experiments. *Journal of Experimental Psychology: Human Perception and Performance*, *7*, 441–458.

Laming, D. R. J. (1997). *The measurement of sensation*. London: Oxford University Press.

Levin, I. R., Johnson, R. D., Russo, C. P., & Delden, P. J. (1985). Framing effects in judgment tasks with varying amounts of information. *Organizational Behavior and Human Decision Processes*, *36*, 362–377.

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

Lópes, L. L., & Oden, G. C. (1983). Some thoughts on the psychological concept of risk. *Journal of Experimental Psychology: Human Perception and Performance*, *9*, 137–144.

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

Mellers, B. A., & Chang, S. (1994). Representations of risk judgments. *Organizational Behavior and Human Decision Processes*, *57*, 167–184.

Mellers, B. A., Chang, S. J., Birnbaum, M. H., & Ordóñez, L. D. (1992). Preferences, prices, and ratings in risky decision-making. *Journal of Experimental Psychology: Human Perception and Performance*, *18*, 347–361.

Mellers, B. A., Ordóñez, L. D., & Birnbaum, M. H. (1992). A change-of-process theory for contextual effects and preference reversals in risky decision making. *Organizational Behavior and Human Decision Processes*, *52*, 331–369.

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

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

Payne, J. W., Bettman, J. R., & Johnson, E. J. (1992). Behavioral decision research: A constructive processing perspective. *Annual Review of Psychology*, *43*, 87–131.

Poulton, E. C. (1989). *Bias in quantifying judgments*. Hillsdale, NJ: Erlbaum.

Poulton, E. C. (1994). *Behavioural decision theory: A new approach*. Cambridge, England: Cambridge University Press.

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

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

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

Sedlmeier, P., & Betsch, T. (Eds.). (2002). *Etc. Frequency processing and cognition*. Oxford, England: Oxford University Press.

Shanteau, J. (1975). Averaging versus multiplying combination rules of inference judgment. *Acta Psychologica*, *39*, 83–89.

Simonson, I. (1989). Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*, *16*, 158–174.

Slovic, P. (1995). The construction of preference. *American Psychologist*, *50*, 364–371.

Sprenger, A., & Dougherty, M. R. (2006). Differences between probability and frequency judgments: The role of individual differences in working memory capacity. *Organizational Behavior and Human Decision Processes*, *99*, 202–211.

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., & Brown, G. D. A. (2004). Sequence effects in categorizing tones varying in frequency. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 416–430.

Stewart, N., Brown, G. D. A., & Chater, N. (2002). Sequence effects in categorization of simple perceptual stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *28*, 3–11.

Stewart, N., Brown, G. D. A., & Chater, N. (2005). Absolute identification by relative judgment. *Psychological Review*, *112*, 881–911.

Stewart, N., Chater, N., & Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, *53*, 1–26.

Stewart, N., Chater, N., Stott, H. P., & Reimers, S. (2003). Prospect relativity: How choice options influence decision under risk. *Journal of Experimental Psychology: General*, *132*, 23–46.

Stewart, N., & Morin, C. (2007). Dissimilarity is used as evidence of category membership in multidimensional perceptual categorisation: A test of the similarity-dissimilarity generalised context model. *Quarterly Journal of Experimental Psychology*, *60*, 1337–1346.

Stewart, N., & Reimers, S. (2008a). *How the distribution of attribute values influences decision under risk*. Unpublished manuscript.

Stewart, N., & Reimers, S. (2008b). *Revealed psychoeconomic functions are not stable*. Unpublished manuscript.

Stewart, N., & Simpson, K. (2008). A decision-by-sampling account of decision under risk. In N. Chater & M. Oaksford (Eds.), *The probabilistic mind: Prospects for Bayesian cognitive science* (pp. 261–276). Oxford, England: Oxford University Press.

Stott, H. P. (2006). Cumulative prospect theory's functional menagerie. *Journal of Risk and Uncertainty*, *32*, 101–130.

Thomas, R. P., Dougherty, M. R. P., Sprenger, A. M., & Harbison, J. I. (2008). Diagnostic hypothesis generation and human judgment. *Psychological Review*, *115*, 155–185.

Tversky, A. (1969). Intransitivity of preferences. *Psychological Review*, *76*, 31–48.

Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323.

Tversky, A., & Koehler, D. J. (1994). Support theory: A nonextensional representation of subjective probability. *Psychological Review*, *101*, 547–567.

Vlaev, I., Chater, N., & Stewart, N. (2007a). Financial prospect relativity: Context effects in financial decision-making under risk. *Journal of Behavioral Decision Making*, *20*, 273–304.

Vlaev, I., Chater, N., & Stewart, N. (2007b). Relativistic financial decisions: Context effects on retirement saving and investment risk preferences. *Judgment and Decision Making*, *2*, 292–311.

von Neumann, M., & Morgenstern, O. (1947). *Theory of games and economic behavior* (2nd ed.). Princeton, NJ: Princeton University Press.

Wedell, D. H. (1991). Distinguishing among models of contextually induced preference reversals. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *17*, 767–778.

Windschitl, P. D., & Chambers, J. R. (2004). The dud-alternative effect in likelihood judgment. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *30*, 198–215.

Windschitl, P. D., & Wells, G. L. (1998). The alternative-outcomes effect. *Journal of Personality and Social Psychology*, *75*, 1441–1423.