



Munich Personal RePEc Archive

Boundedly Rational Expected Utility Theory

Navarro-Martinez, Daniel and Loomes, Graham and Isoni, Andrea and Butler, David and Alaoui, Larbi

Universitat Pompeu Fabra, Warwick Business School, Warwick Business School, Mudoch University, Universitat Pompeu Fabra

15 June 2017

Online at <https://mpra.ub.uni-muenchen.de/79893/>
MPRA Paper No. 79893, posted 01 Aug 2017 05:19 UTC

Boundedly Rational Expected Utility Theory

Daniel Navarro-Martinez

Department of Economics and Business, Pompeu Fabra University, and Barcelona Graduate School of Economics, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain.

Corresponding author. Tel.: +34 935421467. E-mail address: daniel.navarro@upf.edu.

Graham Loomes

Warwick Business School, University of Warwick, Coventry CV4 7AL, UK.

Andrea Isoni

Warwick Business School, University of Warwick, Coventry CV4 7AL, UK.

David Butler

Murdoch School of Management and Governance, Murdoch University, Perth, Australia.

Larbi Alaoui

Department of Economics and Business, Pompeu Fabra University, and Barcelona Graduate School of Economics, Ramon Trias Fargas 25-27, 08005 Barcelona, Spain.

Boundedly Rational Expected Utility Theory

Abstract

We build a satisficing model of probabilistic choice under risk which embeds Expected Utility Theory (EUT) into a boundedly rational deliberation process. The decision maker accumulates evidence for and against alternative options by repeatedly sampling from her underlying set of EU preferences until the evidence favouring one option satisfies her desired level of confidence. Notwithstanding its EUT core, the model produces patterns of behaviour that violate standard axioms, while at the same time capturing the systematic relationship between choice probabilities, response times and confidence judgments, which is beyond the scope of theories that do not take deliberation into account. [92 words]

Keywords: Expected utility; bounded rationality; deliberation; probabilistic choice; confidence; response times.

JEL Classification: D03; D81.

1. Introduction

Economics is often said to be the study of the allocation of scarce resources, of how human beings decide to combine their time and skills with physical resources to produce, distribute and consume. However, economic models may sometimes ignore the fact that arriving at decisions is itself an economic activity and that the hardware and software involved – that is, the human brain and its mental processes – are themselves subject to constraints. Herbert Simon emphasised this point in his 1978 Richard T. Ely lecture, in which he discussed the implications of attention being a scarce resource.¹ In a world where there are many (often complex) choices to be made, spending time on any one decision entails an opportunity cost in terms of the potential fruits of other decisions that might have been considered instead. Being unable to devote unlimited time and attention to every decision they encounter, humans generally have to *satisfice* rather than optimise.

Simon bemoaned the lack of interest among economists in the processes that individuals use when deciding how to allocate their scarce mental resources, and he advocated “building a theory of procedural rationality to complement existing theories of substantive rationality”. He suggested that “some elements of such a theory can be borrowed from the neighboring disciplines of operations research, artificial intelligence, and cognitive psychology”, but noted that “an enormous job remains to be done to extend this work and to apply it to specifically economic problems” (Simon, 1978, pp.14–15).

Although there have been some developments along these lines (e.g., Gilboa and Schmeidler, 1995; Rubinstein, 1998; Gigerenzer and Selten, 2001), the decades that followed Simon’s lecture saw the mainstream modelling of individual decision making – especially with respect to choice under risk and uncertainty – take a different direction. Stimulated by experimental data that appeared to violate basic axioms of rational choice, a number of models appeared at the end of the 1970s and in the early 1980s that sought to provide behavioural alternatives to standard Expected Utility Theory (EUT) – see Starmer (2000) for a review. Typically, these were deterministic models that relaxed a particular axiom and/or incorporated various additional features – e.g., reference points, loss aversion, probability weighting, regret, disappointment – to try to account for certain regularities in observed decisions. While such models provided more elaborate descriptive theories of choice, little or no consideration was given to the mental constraints referred to by Simon. His invocation to build boundedly rational procedural models largely fell by the wayside in the field of risky decision making.

Thus we now have an impressive array of alternative models, each of which can claim to accommodate some (but not all) of the observed departures from EUT. However, being typically deterministic and for the most part process-free, these models have no intrinsic explanation for at least three other pervasive empirical regularities in the data, which may arise from features of the decision-making

¹ This was the year in which Simon was awarded the Nobel Prize in economic sciences “for his pioneering research into the decision-making process within economic organizations” (www.nobelprize.org/nobel_prizes/economic-sciences/laureates/1978/simon-bio.html), and this lecture drew on ideas he had been developing since the 1950s.

process: first, the probabilistic nature of most people's decisions;² second, the systematic variability in the time it takes an individual to respond to different decision tasks of comparable complexity;³ and third, the degree of confidence decision makers (DMs) express about their decisions.⁴

So, in this paper, we propose to explore the direction Simon advocated and investigate the potential for applying a boundedly rational deliberative process to the 'industry standard' model of decision making under risk and uncertainty, EUT. We start by identifying in general terms what is required of a procedural model of preferential choice. We then consider how the various components of such a model might be specified in ways that are in keeping with conventional economic assumptions while at the same time allowing for scarcity of time and attention. The resulting model – which we call Boundedly Rational Expected Utility Theory (BREUT) – generates a number of implications, not just about choice probabilities, but also about process measures such as response times and confidence in the decisions made. A striking result is that, despite being based on EU preferences, the model produces choice patterns that deviate from EUT in line with some well-known decision making phenomena, and we characterise the circumstances under which this occurs.

We shall not claim that the particular mechanisms used in BREUT to model deliberation are a literal representation of the way the mind actually operates. Psychology and neuroscience cannot yet tell us precisely how the brain generates decisions, so BREUT is necessarily a stylisation intended to capture some broad procedural features using conventional modelling elements. Also, while we take EUT as the 'core' of our model, we wish to stress that we are not asserting that EUT is necessarily the correct account of the way people's underlying preferences are configured. As we shall explain in due course, our broad modelling strategy has the potential to be extended to many non-EU core theories. This paper may be understood as a 'proof of concept' exercise: that is, an exploration of the implications of embedding EUT in a simple and reasonable boundedly rational procedure for arriving at decisions.

In the next section, we present our instantiation of BREUT, focusing upon the kind of binary choices between lotteries with monetary outcomes that have been the staple diet of many decision-making experiments. In Section 3, we demonstrate how BREUT provides a parsimonious account of the systematic relationship between choice probabilities, decision time and confidence. We show that the model entails respect for first order stochastic dominance and weak stochastic transitivity but allows patterns of choice that violate strong stochastic transitivity, independence and betweenness. In the final section, we consider the

² When participants are presented with exactly the same decision task displayed in exactly the same way on two or more separate occasions within the same experiment, they often give different responses on different occasions. See Mosteller and Nogee (1951) for an early example; Luce and Suppes (1965) for a review of the early theoretical literature; and Rieskamp *et al.* (2006), Bardsley *et al.* (2010, Chapter 7), and Baucells and Villasís (2010) for more recent discussions. The same pattern of variability has also been observed in more applied settings, such as the judgments of software developers or of medical doctors (see Kahneman *et al.*, 2016).

³ See, for example, Tyebjee (1979), Birnbaum and Jou (1990), Busemeyer and Townsend (1993), Wilcox (1993), Moffatt (2005), Rubinstein (2007, 2013), Spiliopoulos and Ortmann (2016), Achtziger and Alós-Ferrer (2014).

⁴ See, for example, Butler and Loomes (1988), Sieck and Yates (1997), Pleskac and Busemeyer (2010).

relationship between our model and others in the psychology and economics literature. We discuss some limitations of the model in its current form, together with what we see as the most promising directions for extending our approach and some other possible implications of our findings. The Appendix reports the proofs of our theorems.

2. The Model

The term ‘bounded rationality’ has been interpreted in somewhat different ways – see Munier *et al.* (1999) and Harstad and Selten (2013) for examples – so it may be helpful to start by clarifying the way in which we are using it.

Bounded rationality has often been characterised in terms of the difficulties DMs may encounter when they are faced with complex problems or environments involving information that is hard to obtain and process. However, we suggest that even in cases in which there are just two alternatives whose properties are clearly specified, bounded rationality plays a role in terms of the deliberative process used by a DM to arrive at a decision. Straightforward though a choice might appear, it will often still require some thought to identify and evaluate the arguments pulling in opposing directions. Since a DM has many other things to do in her lifetime, it is impossible to devote unlimited time and attention to any single decision. Instead, she has to find a way to allocate her limited attention to the various decisions that she faces – which is to say, she must have some mechanism to decide when to cease further deliberation, make a decision and move on to the next task. It is the process underlying the allocation of time and attention between different decisions that we regard as fundamental to the characterisation of bounded rationality.

Behavioural scientists have invested substantial effort in developing various approaches for modelling decision-making processes: see, for example, elimination by aspects (Tversky, 1972); the adaptive decision maker framework (Payne *et al.*, 1993); the priority heuristic (Brandstätter *et al.*, 2006); and query theory (Johnson *et al.*, 2007), to name only a few. Another influential stream of literature has developed an *accumulator* or *sequential-sampling* framework (for reviews, see Ratcliff and Smith, 2004; Otter *et al.*, 2008). Most such accumulator models have been developed in the context of perceptual tasks (e.g., judging the colour or comparative brightness of objects, or the relative numbers of dots in different areas of a screen, or the relative lengths of lines, and so on). The application of these models to preferential choice has been less common, although Busemeyer and Townsend’s (1993) Decision Field Theory (DFT) and its extensions constitute a notable exception. We will discuss the relationship between BREUT and DFT in more detail in Section 4.

The general idea behind such models is that a DM starts from a position where she does not already have a decisive preference for any one of the available options and therefore has to acquire evidence that can help to discriminate between them. In the case of preferential choice, this evidence comes from introspecting about the relative (subjective) desirability of the possible consequences and the weights to be attached to each of them. As evidence is accumulated, it is as if the DM continually re-assesses the relative strengths of the arguments for each option. As long as no single option is favoured sufficiently strongly over the other(s),

further sampling of evidence – that is, further deliberation – is required. However, once the judged relative advantage of one option over the other(s) crosses some threshold, the process is terminated and the option favoured by the accumulated evidence is chosen. The DM can then allocate her scarce attention to something else.

Such models typically consist of three components:

- (i) some representation of the sources of evidence from which samples are drawn – in this context, the underlying stock of subjective values or judgments;
- (ii) some account of the way the sampled evidence is accumulated;
- (iii) a stopping rule which terminates the accumulation process and triggers a decision.

In the next three subsections, we explain how BREUT models these three components and we describe the resulting behavioural variables.

2.1. The Structure of Underlying Subjective Values

Consider the following seemingly simple decision. A DM is asked to choose between lottery A , which (omitting any currency symbol) offers the certainty of 30, and lottery B , which offers a 0.8 chance of 40 and a 0.2 chance of 0. She might be told that the uncertainty will be resolved by drawing a ticket at random from a bag containing 100 tickets numbered from 1 to 100 inclusive: if she chooses A she will get 30 whichever ticket is drawn; if she chooses B she will receive 40 if the ticket is between 1 and 80 inclusive, while a ticket between 81 and 100 gives 0.

Using standard EUT notation, with \succcurlyeq denoting weak preference, the decision can be written as:

$$A \succcurlyeq B \Leftrightarrow u(30) \geq 0.8 u(40) + 0.2 u(0). \quad (1)$$

Deterministic EUT supposes that it is as if each DM acts according to a single utility function that gives an exact answer to this question (and gives the same answer every time that this question is asked). But cognitive psychology and neuroscience suggest that there is no unique and instantly accessible subjective value function – see, for example, Busemeyer and Townsend (1993), Gold and Shadlen (2007), Stewart *et al.* (2006, 2015). Rather, a typical DM will have had many experiences and impressions of what 30 and 40 represent in subjective terms. So it may require some retrieval from memory and some reflection to balance the arguments pulling in opposite directions. On the one hand, lottery B offers a 0.8 chance of a higher payoff – 40 rather than 30; but A offers a 0.2 chance of getting 30 rather than 0. For a typical DM, it may not be immediately obvious exactly where the subjective value of 30 is located in the range between the subjective values of 0 and 40 nor precisely how the differences, weighted by the probabilities, balance out. So arriving at a decision may involve deliberating about the balance of evidence obtained by sampling from those experiences and judgments.

Since we are investigating the effects of embedding EUT in a boundedly rational deliberation process, we represent the underlying stock, or *core*, of past experiences and impressions by a distribution of

von Neumann-Morgenstern (vNM) utility functions, $u(\cdot)$, normalised so that $u(0) = 0$.⁵ At this point, we place no restrictions upon the nature of this distribution, although in Section 3 we make more specific assumptions in order to explore the behaviour of the model and derive predictions for commonly used pairs of lotteries.

The assumption of a distribution of vNM functions was made by Becker *et al.* (1963a) when they proposed their “random utility model for wagers”, which has since come to be known as the *random preference* (RP) model. This way of representing the variability from moment to moment of people’s states of mind might be seen as an early example of a ‘multiple-selves’ modelling strategy that has been applied to various areas of economic behaviour (e.g., Elster, 1987; Alós-Ferrer and Strack, 2014). The key difference between the RP model and ours is that Becker *et al.* supposed that DMs only sampled once for each decision, applying a single randomly-picked $u(\cdot)$ to both options, whereas we assume that deliberation involves sampling multiple times, as we now describe.

2.2. Modelling the Sampling and Accumulation of Evidence

Adapting the general framework of accumulator models to the specific context of binary choice between lotteries, in BREUT each sample is modelled as a random draw of a utility function from the core distribution of $u(\cdot)$, which is applied to the pair of lotteries under consideration. Using probabilities as weights, as EUT entails, this yields a subjective value difference which we denote by $V(A, B)$ and which takes some positive value when a $u(\cdot)$ is drawn that strictly favours A , takes a value of zero when the two options are exactly balanced and takes a negative value when the sampled $u(\cdot)$ strictly favours B . We can represent this difference as a difference between the monetary certainty equivalents (CEs) of the two options. Formally, for any $u(\cdot)$ sampled,

$$V(A, B) = CE_A - CE_B = u^{-1}(EU_A) - u^{-1}(EU_B). \quad (2)$$

We use differences in CEs rather than in utilities because CEs are measures that can be legitimately compared and aggregated across utility functions. It is well known in economics that comparisons of utilities (or their aggregation) across different utility functions are not theoretically meaningful and lead to problematic results. This is true even if utilities are normalized, for example to be between 0 – for the worst possible outcome in a particular set – and 1 – for the best possible outcome (see Hammond, 1993; Binmore,

⁵ In EUT, as presented by von Neumann and Morgenstern (1947), the argument(s) of $u(\cdot)$ could be anything – money, consumption goods, health states, etc. In most applications to experimental data involving monetary lotteries (e.g., Hey and Orme, 1994; Harless and Camerer, 1994; Holt and Laury, 2002; Blavatskyy and Pogrebna, 2010), the argument is taken to be the outcome of the lottery *as described to the participant* – i.e., ignoring any other wealth they might have access to. Given the relevance of what we do for this literature, we also follow that approach. As shown by Rabin (2000), assuming that decision makers integrate the consequences of the lotteries they face with their (lifetime) wealth, may lead to paradoxical implications; but whether that assumption is an integral part of EUT is contested (e.g., Rubinstein, 2001; Cox and Sadiraj, 2006; Palacios-Huerta and Serrano, 2006; Rubinstein, 2016).

2009, for a detailed discussion of this topic). This argument is typically addressed to comparison or aggregation of utilities across people, but the same logic holds for the case of having multiple utility functions that represent different mental states within one person (as if there were multiple selves). There is no basis for claiming that the difference between the best and the worst possible outcomes represents the same amount of satisfaction in one mental state (or self) as in another one. For this reason, we need to put $V(A, B)$ in a ‘common currency’ that can be legitimately aggregated. The use of *CEs* is a straightforward way of achieving that, which has a number of precedents in the literature (e.g., Luce, 1992; Luce *et al.*, 1993; Cerreia-Vioglio *et al.*, 2015).

For each sampled $u(\cdot)$, the *CE* difference provides a signal not only about which option is better, but also about how much better it is. Repeated sampling (with replacement) produces a series of realisations of $V(A, B)$, which are accumulated by progressively updating their mean and sample standard deviation, denoted respectively by $\overline{V(A, B)}$ and $s_{V(A, B)}$.

As noted at the end of Section 1, what we are proposing is a model, not a literal account of how the human brain actually works. We are not suggesting that people really precisely calculate $V(A, B)$ at every moment. Rather, we are exploring a way of modelling aggregation of sampled differences which is consonant with a deliberative process applied to an EUT core, and which could also be extended to other core theories that researchers might wish to consider in future work, given that most decision theories entail the existence of *CE* differences between pairs of alternatives.

This way of aggregating evidence has important implications. Let $E[V(A, B)]$ denote the mean of the distribution of *CE* differences for the pair of lotteries $\{A, B\}$ implied by the individual’s underlying population of $u(\cdot)$. If the individual takes a sample of size k from this distribution, that sample will have a mean which we denote by $\overline{V_k(A, B)}$. Taking sufficiently many samples of size k will result in some distribution for $\overline{V_k(A, B)}$. Independently of the value of k , this distribution will always have a mean equal to $E[V(A, B)]$. However, as implied by the central limit theorem, when k becomes larger, the distribution of $\overline{V_k(A, B)}$ will be increasingly similar to a normal distribution with a variance inversely related to k . If the individual were to deliberate indefinitely – that is, if k were allowed to tend towards infinity – the variance of the distribution of $\overline{V_k(A, B)}$ would tend towards zero and the whole distribution would collapse towards $E[V(A, B)]$. This straightforward implication of the central limit theorem will be particularly useful to shed light on the operation of our model in the following sections. In essence, if an individual had unlimited time to devote to any one decision, that would constitute a *limiting case* in which she would always arrive at the same judgment about the sign and size of $\overline{V(A, B)}$, as given by $E[V(A, B)]$, with no variability. However, in reality, unlimited deliberation is infeasible, so the individual will need to decide when the accumulated evidence is sufficient to make a decision. This is the role of the stopping rule.

2.3. Modelling the Stopping Rule

In a number of models, including Busemeyer and Townsend's (1993) DFT, the process of accumulating evidence in a binary choice can be represented as a random walk or a diffusion process which terminates when the path crosses a (typically fixed) threshold. However, it is often unclear what determines the position of the thresholds, and it is not obvious that they should be fixed. On the one hand, for decisions in which there is a lot at stake and for which there is considerable variance in the evidence, we might expect the thresholds to be set a relatively long way from zero in order to avoid jumping too quickly to what might be the 'wrong' choice. But until the DM has collected some evidence, she has little information about the variance and therefore cannot set the initial threshold distance appropriately. On the other hand, if there is actually little difference between the options, then having fixed parallel thresholds entails the possibility that it could take an unrealistically long time to reach one or other threshold in cases in which the consequences of picking the wrong option are negligible.

To address these issues, we propose an approach which, in effect, sets thresholds that are responsive both to the evolving pattern of the evidence as it accumulates and to the DM's wish to limit the time spent deliberating about any particular decision. The key to our stopping rule is the DM's *desired level of confidence*: in particular, we propose that the DM deliberates until she concludes that the accumulated evidence gives her sufficient confidence to make a choice. In BREUT, confidence is represented as the probability that the DM picks the option that she would choose after unlimited deliberation – i.e., the option implied by the sign of $E[V(A, B)]$.⁶

We suggest that the notion of individuals attempting to achieve a personal desired level of confidence is a simple and intuitive way of building a satisficing model, and also one for which there is some empirical support (see Hausmann and Läge, 2008). In contrast with the optimal stopping tradition (see, e.g., Stigler, 1961; Shiryaev, 1978), this approach does not require us to assume that the individual has detailed knowledge about the opportunity cost of additional sampling in terms of forgone benefits from potential future activities. We simply require that the DM can form an impression of the confidence level she wants at the time the decision is made.⁷

We denote the DM's desired level of confidence as *Conf*. Because deliberation is costly in terms of the opportunity costs associated with each extra draw, we allow the DM to progressively reduce *Conf* as the amount of sampling increases. The idea here is that the longer she spends trying to discriminate between options, the more likely she is to conclude that there is not much between them, so that she has less to fear from choosing the wrong option. Specifically, we assume that the desired level of confidence after k draws is given by:

⁶ If $E[V(A, B)] = 0$, the DM is indifferent between A and B in the limit. We shall assume that, in such a case, she chooses each option with probability 0.5.

⁷ It will be clear from the following sections that our main choice results do not depend on the use of this particular stopping rule.

$$Conf = \max [0.5, 1 - d(k - 1)], \quad (3)$$

where $k \geq 2$ and where d (with $0 < d \leq 0.5$) is a parameter that captures the rate at which the DM reduces her $Conf$ as k increases, subject to the constraint that $Conf \geq 0.5$.

d may vary from one individual to another, reflecting different tastes for the trade-off between more input into the current decision and turning attention to something else. A person with a very low d is someone who wants to be very confident in her decisions, and therefore is willing to invest more time deliberating. The limiting case is when $d \rightarrow 0$, in which the individual wants to be absolutely sure of making the right decision and deliberates indefinitely. On the other hand, someone with a high value of d is ready to make decisions with less confidence and spends relatively less time deliberating. Thus when d has the maximum value of 0.5, the DM makes her decision after just two samples, choosing the option favoured by the mean of the two sampled CE differences. Modelling d as a personal characteristic provides a degree of within-person consistency with just one parameter. It would be possible to construct a more complicated function for $Conf$, but the linear form in Expression 3 is sufficient for our ‘proof of concept’ purposes.⁸

We then model the DM’s decision about whether or not she should terminate her deliberation as if she were applying a sequential t -test. Again, we acknowledge the *as if* nature of this assumption. Other ways of modelling the stopping rule are possible, but a sequential statistical test meshes well with the idea of achieving some personal level of confidence. What we try to capture is the idea that when the choice is initially presented – i.e., before any deliberation has occurred – the DM starts with the null hypothesis that there is no significant difference between the subjective values of the two options. However, as the evidence accumulates, it is as if she continually updates $\overline{V(A, B)}$ and $s_{V(A, B)}$ and combines them to form a test statistic T_k :

$$T_k = \frac{\overline{V(A, B)}}{s_{V(A, B)}/\sqrt{k}}. \quad (4)$$

This statistic is then used to determine whether the null hypothesis of zero difference can be rejected at the level of $Conf$ corresponding to k . This occurs if the following condition is met:

$$F_{k-1}[\text{abs}(T_k)] \geq Conf, \quad (5)$$

where $F_{k-1}[\cdot]$ is the c.d.f. of the t -distribution with $k - 1$ degrees of freedom. If the weak inequality in Expression 5 is not satisfied, the DM is assumed to continue sampling and to continue reducing $Conf$ until the hypothesis of zero difference is rejected in favour of one of the alternatives – at which point, she chooses whichever option is favoured according to the sign of $\overline{V(A, B)}$. The value of k when sampling stops and a

⁸ For instance, one could allow the DM to reduce her desired level of confidence more slowly when the stakes are higher, or one could assume some nonlinear relationship with k .

choice is made is denoted by k^* , the level of confidence at that point is $Conf^*$ and the value of the test statistic at that point is T_{k^*} .

2.4. Behavioural Variables Generated by BREUT

Because of its procedural nature, BREUT gives a richer description of decision making than process-free models. The outcome of the deliberation process can be summarised by three main behavioural variables: choice probabilities, confidence and response times. We now consider each of these in more detail.

2.4.1. Choice probabilities

In BREUT, the probability of choosing A over B , denoted by $\Pr(A \succ B)$, is the probability that the null hypothesis is rejected with a positive T_{k^*} . The complementary probability that B is chosen is the probability that the null hypothesis is rejected with a negative T_{k^*} .

In the classic RP model, it is as if an individual samples just once per decision and chooses on the basis of the single $u(\cdot)$ sampled on that occasion. In that one-shot model, the probability of choosing A over B is given by the proportion of utility functions that favour A over B in the underlying core distribution. We denote this probability by $\Pr Core(A \succ B)$. However, BREUT supposes that an individual samples more than once, so that the effective $\Pr(A \succ B)$ will typically be different from $\Pr Core(A \succ B)$. As noted in Section 2.2, the central limit theorem implies that the variance of the distribution of $\overline{V_k(A, B)}$ decreases as k becomes larger and the whole distribution gets more concentrated around $E[V(A, B)]$. As a consequence, as k increases, the DM becomes increasingly likely to choose the option favoured by the sign of $E[V(A, B)]$. In the limit, the probability of choosing A over B as k tends towards infinity, denoted by $\Pr Lim(A \succ B)$, is either 0 (if $E[V(A, B)]$ is negative) or 1 (if $E[V(A, B)]$ is positive). This means that:

$$\begin{aligned} 0 = \Pr Lim(A \succ B) \leq \Pr(A \succ B) \leq \Pr Core(A \succ B) & \quad \text{if } E[V(A, B)] < 0 \\ \Pr Core(A \succ B) \leq \Pr(A \succ B) \leq \Pr Lim(A \succ B) = 1 & \quad \text{if } E[V(A, B)] > 0. \end{aligned} \tag{6}$$

That is, $\Pr(A \succ B)$ will always lie between the core probability in the one-shot RP model and the limiting probability (0 or 1) implied by the sign of the mean of the distribution of CE differences.

Expression 6 has some important implications. First, if one of the lotteries is never favoured by the individual's core utility functions (i.e., $\Pr Core(A \succ B)$ is either 0 or 1), then there will be no amount of sampling for which that lottery will be chosen with positive probability (i.e., it will always be the case that $\Pr(A \succ B) = \Pr Core(A \succ B) = \Pr Lim(A \succ B)$). Since dominated lotteries are never chosen in EUT, this entails that BREUT necessarily satisfies first-order stochastic dominance. Second, since (other things being equal) lower values of d imply larger values of k^* , $\Pr(A \succ B)$ will tend towards $\Pr Lim(A \succ B)$ as d decreases: in other words, with more deliberation, choice probabilities become more extreme. Third, if $\Pr Lim(A \succ B)$ and $\Pr Core(A \succ B)$ are on different sides of 0.5 for some core, then BREUT allows for one lottery to be the modal choice even though the majority of the DM's core utility functions favour the other lottery. Theorems

1 and 3 in Section 3 will show how this allows a stable core of vNM utility functions to produce Allais-type violations of EUT's independence axiom.

2.4.2. Confidence

In BREUT, the degree of confidence in a decision, $Conf^*$, is the level of $Conf$ used in the test that rejects the null hypothesis and triggers a choice in favour of one of the alternatives after k^* samples. The range of the confidence measure is $(0.5, 1)$ – see Expression 3.

While a substantial literature has investigated and modelled confidence judgments in decision tasks that have a correct answer (see Pleskac and Busemeyer, 2010, and references therein), there are only a few isolated exceptions of studies investigating confidence in preferential choice (e.g., Butler and Loomes, 1988; Sieck and Yates, 1997). To the best of our knowledge, there are no other decision models that explicitly address confidence in preferential decisions. However, confidence is clearly regarded as an important factor in various areas of economics (e.g., Acemoglu and Scott, 1994; Ludvigson, 2004; Dominitz and Manski, 2004; Barsky and Sims, 2012) and it might be desirable if decision theories had more to say about this dimension. A distinctive feature of BREUT is that its stopping rule automatically produces predictions about the DM's confidence in the decision made.⁹

2.4.3. Response time

Since deliberation is modelled through sequential sampling, BREUT can make predictions about the length of time that it takes to make a decision – the *response time* (RT). It is reasonable to assume that RT is an increasing function of the number of samples, k^* , required to reject the null hypothesis. In addition, RT can also be expected to be positively related to the complexity of the decision problem. As there is no generally agreed index of complexity for such choices, for the purposes of this paper we make the simple assumption that complexity is reflected by the total number of consequences (NC) appearing in any pair.¹⁰ This is a straightforward way of capturing the intuition that, if there are more items to consider, each deliberation step will take longer. So we propose the following basic specification for RT :

$$RT = k^* \cdot NC. \quad (7)$$

Since k^* is a stochastic variable, RT is also stochastic, implying that the DM is liable to take different amounts of time to reach a decision for the same pair of alternatives on different occasions. Given Expression 3 and the definition of k^* , it follows that $RT = [(1 - Conf^*)/d + 1]NC$. That is, RT and $Conf^*$ are

⁹ There may also be interesting questions to be asked about the effects on behaviour when DMs are prevented from achieving their desired level of confidence – for example, when operating under time constraints (e.g., Kocher *et al.*, 2013) and having to make a choice before the null hypothesis can be rejected at the unconstrained desired level of confidence. But these questions go beyond the scope of the present paper.

¹⁰ If the same consequence appears in both alternatives, it is counted twice in NC .

inversely related. Also, RT only depends on the other parameters of the model and is not directly affected by any other contextual features. In practice, the magnitude of response times is likely to be affected by other factors, such as the particular ways in which the alternatives are displayed, the time needed to scan the stimuli, the respondent's familiarity with the nature and format of the task, her degree of fatigue and so on. More elaborate expressions could be constructed to allow for such considerations, including a scaling factor that maps RTs to real time if the model is to be fitted to data, but Expression 7 is adequate for our immediate purposes.

3. Exploring the predictions of BREUT

This section is devoted to the exploration of the predictions of BREUT. Here we will illustrate how systematic changes in the core distribution and the free parameter d affect the behavioural variables described in Section 2.4, and we will apply BREUT to particular decision problems that have played a major role in motivating the development of non-EU models.

Unlike many deterministic models, more complex, dynamic stochastic processes such as the ones we are considering do not necessarily lend themselves to precise analytical results. For this reason, we use an approach based on a combination of theoretical and simulation results. Our results for positive values of d (i.e., for limited sampling) are based on simulations.¹¹ These, in combination with Expression 6, provide a clear picture of how the model behaves when its key parameters are systematically varied. For the more specific decision problems that we examine, we also provide theoretical results based on theorems for the limiting case of $d \rightarrow 0$.¹²

Throughout this section, we assume that the DM's underlying subjective values can be represented by a fixed distribution of vNM utility functions, each of which takes the form:

$$U(x) = x^{1-r}, \quad (8)$$

with $r < 1$. This type of utility function has been widely used in the decision-making literature (see Stott, 2006) and has the property of constant relative risk aversion (CRRA). Specifically, with Expression 8, the DM is risk averse for $0 < r < 1$, risk neutral for $r = 0$, and risk seeking for $r < 0$. So, as r increases, the DM becomes more risk averse. Other specifications could be used without fundamentally changing the main implications of the model.¹³

¹¹ Simulation methods are becoming increasingly popular in economics (e.g., Calvó-Armengol and Jackson, 2004; Reiss, 2011; Elliott *et al.*, 2014) and they are widely used in econometrics (see, e.g., Gouriéroux and Monfort, 1997).

¹² In addition to the simulation results presented here, we have conducted various robustness checks that are available upon request.

¹³ We have conducted a variety of simulations using an underlying distribution of CARA utility functions of the form $U(x) = (1 - e^{-rx})/r$, and also simulations restricting all functions to be concave, the results of which are available from the authors upon request.

To simulate BREUT, we also assume that each time the DM samples from her core preferences, it is as if an r is extracted from a transformed beta distribution of risk attitude parameters, such that:

$$r \sim \text{Beta}(3, 3) \cdot \beta + \left(\alpha - \frac{\beta}{2}\right). \quad (9)$$

This means that her r values are drawn from a symmetric and bell-shaped distribution with mean α and range β , which is bounded below and above at $\alpha - \beta/2$ and $\alpha + \beta/2$ respectively. If $\beta = 0$, the specification reduces to deterministic EUT with $r = \alpha$. There are three main advantages of using this specific distribution of r values. First, it is described by only two intuitive parameters, mean and range, that capture location and dispersion. Second, it is symmetric, which makes our results more straightforward and shows that they do not depend on asymmetries in the distribution of risk aversion parameters. Third, the distribution is bounded, which avoids sampling extreme or implausible risk aversion parameters. Other similar distributions could be used without fundamentally affecting the key predictions of the model. Moreover, as explained in Section 3.3, our main theoretical results do not rely on any particular distribution of risk aversion parameters.

Using this core structure, we implement the model by making independent random draws of values of r , each of which entails a *CE* difference between the two lotteries under consideration, and by accumulating those differences according to Expression 4 until the condition in Expression 5 is met. We simulate this process a large number of times for each binary choice (20,000 times unless otherwise stated) to generate predictions about the three behavioural variables described in Section 2.4. The measures reported for confidence and response times are averages across all simulation runs; the choice probabilities are calculated as the proportion of times that the corresponding option is chosen.¹⁴

In the rest of the section, we examine how the model behaves in three respects. First, we show the results of comparing a fixed lottery to a monotonic sequence of sure alternatives when we hold the parameters of the model constant. Second, we explore how changes in the three free parameters – α , β and d – affect BREUT's predictions when we hold the alternatives constant. Third, we use more specific sets of decision problems to illustrate how BREUT behaves in scenarios involving stochastic dominance, transitivity, independence, and betweenness.

¹⁴ All simulations and calculations were implemented using the R statistical programming language. The code is available from the authors on request.

3.1. Fixed Lottery vs Variable Sure Amounts

Table 1 shows how choice probabilities, confidence and response times vary in choices between a fixed lottery and a series of sure amounts of money, when we hold the free parameters constant at $\alpha = 0.35$, $\beta = 1.0$ and $d = 0.1$. In the context of CRRA preferences, these parameters are very much in line with estimates derived from experimental data (e.g., Holt and Laury, 2002). The fixed lottery B (shown in the heading of the table) offers a payoff of 40 with probability 0.8 and zero with probability 0.2, represented as **(40, 0.8; 0, 0.2)**. The sure amounts A (shown in the first column) increase from 20 to 32 in steps of 2.

[Insert Table 1 here]

As should be expected, increasing the value of A raises the proportion of utility functions favouring A over B , so that $\Pr \text{Core}(A \succ B)$ rises from 0.04 (when A is 20) to 0.97 (when A is 32). For the values of A from 20 to 26 inclusive, the mean of the core distribution of CE differences is negative, as shown in the $E[V(A, B)]$ column, so that $\Pr \text{Lim}(A \succ B) = 0$.¹⁵ In these cases, repeated sampling moves $\Pr(A \succ B)$ away from $\Pr \text{Core}(A \succ B)$ towards 0, as entailed by Expression 6. Even when 26% of $u(\cdot)$ favour A , as in the fourth row where $A = \mathbf{(26, 1)}$, repeated sampling results in A being picked on only 13% of occasions.

However, when A increases to **(28, 1)**, $E[V(A, B)]$ becomes positive, so that $\Pr \text{Lim}(A \succ B) = 1$ for this and all higher values of A . This means that for the remaining rows in the table, repeated sampling moves $\Pr(A \succ B)$ above $\Pr \text{Core}(A \succ B)$. As a consequence, even though just 45% of $u(\cdot)$ favour A when $A = \mathbf{(28, 1)}$, A is chosen more often – in this example, in slightly more than 50% of the simulation runs, resulting in $\Pr(A \succ B)$ and $\Pr \text{Core}(A \succ B)$ being on different sides of 0.5. When $A = \mathbf{(30, 1)}$ and is favoured by 73% of $u(\cdot)$, the deliberative process results in A being chosen in 92% of runs, showing again the tendency implied by Expression 6.

The next two columns show the average response times, RT , and the average levels of Conf^* . In the rows towards the top and bottom of the table, where one or the other option is strongly favoured, RT s are shorter and Conf^* is higher. In the middle rows, where the two options are more finely balanced, more sampling is required to reject the null hypothesis of zero difference, so that RT s become longer and Conf^* decreases. This pattern in the RT s is in line with existing empirical evidence for choices between lotteries (see, e.g., Mosteller and Nogee, 1951; Jamieson and Petrusic, 1977; Petrusic and Jamieson, 1978; Moffatt, 2005).¹⁶ The confidence pattern is in line with the evidence found in perceptual tasks (see Pleskac and Busemeyer, 2010, for a review of this literature).

¹⁵ $E[V(A, B)]$ has been calculated in all cases as the mean of 100,000 simulated CE differences.

¹⁶ The overall distribution of RT s tends to be positively skewed. This is often observed in experiments (e.g., Ratcliff and Smith, 2004), even though this evidence typically comes from tasks in which there is a correct answer, as opposed to preferential choice.

3.2. Changing the Free Parameters of the Model

Tables 2, 3 and 4 show how the model behaves for a fixed lottery pair when each of its free parameters (α , β and d) is progressively changed, holding all other parameters constant. For illustrative purposes, the exercise is conducted for lottery pair $\{A, B\}$, with $A = (\mathbf{30}, 1)$ and $B = (\mathbf{40}, 0.8; \mathbf{0}, 0.2)$. The main results can be summarised as follows.

Table 2 shows that when α (the mean value of r) is progressively increased so that the DM becomes more risk averse overall, $\Pr(A \succ B)$ and $\Pr \text{Core}(A \succ B)$ both increase monotonically, with $\Pr(A \succ B)$ tending to be more extreme than $\Pr \text{Core}(A \succ B)$, as implied by Expression 6. RT s tend to increase for more finely balanced decisions, while the Conf^* values show the opposite pattern.

[Insert Table 2 here]

When β is increased, widening the range of the distribution of r and producing greater variability in $V(A, B)$, Table 3 shows that more sampling is required to trigger a decision, entailing longer RT s and lower Conf^* .

[Insert Table 3 here]

Finally, Table 4 shows the effect of decreasing d . When $d = 0.5$ (so that $k^* = 2$), $\Pr(A \succ B)$ goes from the $\Pr \text{Core}(A \succ B)$ level of 0.64 to 0.75 (note that since $E[V(A, B)] > 0$, $\Pr \text{Lim}(A \succ B) = 1$ in this case). At the other extreme, when $d = 0.01$, $\Pr(A \succ B)$ is much closer to the limiting probability of 1. Lower values of d mean that the DM is less willing to reduce her desired level of confidence, which entails increasing the average amount of sampling and the average RT s. So in the third row, where $d = 0.3$, average RT s are 7.01, reflecting the fact that with $NC = 3$, the average k^* is 2.34; in the sixth row, where $d = 0.1$, RT s are around 10. The RT s continue to rise as d falls further.

[Insert Table 4 here]

3.3. Implications for Stochastic Dominance, Transitivity, Independence and Betweenness

We now apply BREUT to some specific problems which are typical of those used in many of the experimental studies that have motivated much of the theory development in decision making under risk and uncertainty. We explore to what extent BREUT's predictions do or do not correspond with various well-known patterns. For each type of scenario, we start from theoretical results for the limiting case and then illustrate the key patterns using simulations.

First, we show how the model behaves in conditions of first order stochastic dominance (FOSD). Next, we consider what scope there is for violations of stochastic transitivity within the BREUT framework. We then turn to violations of the independence axiom and the betweenness property of EUT.

3.3.1. First Order Stochastic Dominance

As noted in Section 2.4, Expression 6 entails that whenever one of the lotteries is never favoured by the core utility functions there is no amount of sampling for which that lottery will be chosen with positive probability. This condition is trivially satisfied in the case of FOSD. Our simulations will look at the behaviour of BREUT's procedural measures, RT and $Conf^*$, for different FOSD lottery pairs.

The two pairs in Table 5 involve transparent FOSD. All lotteries offer 50-50 chances of zero or a positive payoff, with A offering a higher positive payoff than B in both pairs, 10 more in the first and 1 more in the second, so that $\Pr Core(A \succ B)$, $\Pr Lim(A \succ B)$ and $\Pr(A \succ B)$ all equal 1. In spite of the larger payoff difference in favour of A in the first pair, decisions are reached quickly and with high confidence in both cases, resulting in virtually identical RT s and $Conf^*$.¹⁷

[Insert Table 5 here]

3.3.2. Weak and Strong Stochastic Transitivity

In the probabilistic choice literature, a distinction has been made between weak stochastic transitivity (WST) and strong stochastic transitivity (SST). For any three options X, Y, Z , WST requires that if $\Pr(X \succ Y) \geq 0.5$ and $\Pr(Y \succ Z) \geq 0.5$, then $\Pr(X \succ Z) \geq 0.5$. The stronger requirement in SST is that if $\Pr(X \succ Y) \geq 0.5$ and $\Pr(Y \succ Z) \geq 0.5$, then $\Pr(X \succ Z)$ must be at least as large as the greater of those two: $\Pr(X \succ Z) \geq \max[\Pr(X \succ Y), \Pr(Y \succ Z)]$. As Tversky and Russo (1969) showed, SST is equivalent to an *independence between alternatives* condition, whereby $\Pr(X \succ Z) \geq \Pr(Y \succ Z)$ if and only if $\Pr(X \succ W) \geq \Pr(Y \succ W)$ for any X, Y, Z and W . That is, the relationship between the probabilities that each of two lotteries is chosen over a common alternative should not be reversed if the alternative is changed.

Rieskamp *et al.* (2006) concluded that the empirical evidence of violations of WST was thin, whereas there was plentiful evidence of violations of SST. In this subsection we show that BREUT is consistent with WST – i.e., the only instances of violations of WST will be due to random variation rather than to any systematic underlying tendency – but it allows systematic violations of SST of the kinds that have been documented.

In order for there to be any tendency to violate WST in the limit, it would be necessary to generate a case in which the mean CE differences for $\{X, Y\}$ and $\{Y, Z\}$ are positive but the mean CE difference for $\{X, Z\}$ is negative. However, this is clearly impossible in the limiting case, as $E[V(X, Y)] + E[V(Y, Z)] = E[V(X, Z)]$. Also, any set of $u(\cdot)$ will yield a mean CE for each of X, Y and Z , and if $CE(X) > CE(Y)$ and $CE(Y) >$

¹⁷ Empirically, even when FOSD is transparent, there are occasional violations, and there is no discernible difference in this respect between pairs in which the increment is larger or smaller (see Butler *et al.*, 2014a, 2014b). One possibility would be to ascribe these to pure mistakes (e.g., simple lapses of concentration) and capture them by a *tremble* term, as in Moffatt and Peters (2001) or Loomes *et al.* (2002). Substantial rates of violation that cannot be ascribed to such factors (e.g., the violations that Birnbaum, 2008, has explained using his TAX model) could be interpreted as evidence against the assumption of an EUT core.

$CE(Z)$, then $CE(X) > CE(Z)$. Thus all three differences must have the same sign, so that lower values of d and longer RT s will tend to push $\Pr(X \succ Y)$, $\Pr(Y \succ Z)$ and $\Pr(X \succ Z)$ progressively closer to 1.¹⁸

However, violations of SST are a different matter. This can be conveniently illustrated with an example of the so-called *Myers effect* (Myers *et al.*, 1965), which constitutes a violation of the independence condition as specified by Tversky and Russo (1969). Table 8 shows how the probabilities of choosing between each of two lotteries K and L and a series of sure amounts (M) change as the sure sums are progressively increased. The independence condition applied to this case entails that any inequality between $\Pr(K \succ M)$ and $\Pr(L \succ M)$ should hold for all values of M . As Table 6 shows, this is not the case.

[Insert Table 6 here]

Because K has a wider range of payoffs than L , as the sure amount M is increased, $\Pr \text{Core}(K \succ M)$ in the top panel of Table 6 changes more slowly than $\Pr \text{Core}(L \succ M)$ in the bottom panel. For values of M below 30, $\Pr \text{Core}(K \succ M) < \Pr \text{Core}(L \succ M)$, while the opposite is true for M above 30. That is, violations of SST are already present in the RP core prior to any sampling. As implied by Expression 6, the tendencies in core probabilities are amplified by repeated sampling. Since for both lotteries the mean CE difference changes sign as M is increased, BREUT produces even more marked violations of SST.

The patterns of RT and $Conf^*$ are again closely related to patterns in choice probabilities. In the top panel, choice probabilities vary less over the range of M that we consider, and RT and $Conf^*$ also display limited variation. There is more variation in choice probabilities in the bottom panel, matched by more pronounced increases and decreases in RT s and opposite patterns in $Conf^*$. As seen earlier, when probabilities are more extreme (as in the bottom panel), decisions are taken more quickly and with greater confidence than when they are closer to 0.5 (as in the bottom panel). Given the similar level of complexity of the pairs involving K and L , when $Conf^*$ in the bottom panel is close to the level of the top panel, RT values are also similar.

3.3.3. Implications for Independence and Betweenness

Many experimental tests of the independence axiom of deterministic EUT have used pairs of lotteries that can be represented in Marschak-Machina (M-M) diagrams such as that shown in Figure 1 (e.g., Marschak, 1950; Machina, 1982). For any three distinct money payoffs, $x_h > x_m > x_l \geq 0$, the vertical axis shows the probability of receiving x_h and the horizontal axis shows the probability of being paid x_l , with any residual probability assigned to x_m . Figure 1 has been drawn for $x_h = 40$, $x_m = 30$ and $x_l = 0$, with $A = (30, 1)$, $B = (40, 0.8; 0, 0.2)$, $C = (30, 0.25; 0, 0.75)$, $D = (40, 0.2; 0, 0.8)$ and $E = (40, 0.2; 30, 0.75; 0, 0.05)$.

[Insert Figure 1 here]

¹⁸ It may be possible to produce sets of $u(\cdot)$ that give violations of WST under RP's single-sample conditions, but in BREUT such patterns will be attenuated by repeated sampling.

In any such triangle, deterministic EUT entails that a DM's preferences can be represented by a set of linear and parallel indifference curves sloping up from south-west to north-east, with the gradient of the lines reflecting her attitude to risk (see Machina, 1982). The straight and parallel nature of the indifference curves entails that an EU maximiser who chooses A (respectively, B) from pair $\{A, B\}$ in Figure 1, or A (E) from pair $\{A, E\}$, would also choose C (D) from pair $\{C, D\}$. This is an implication of EUT's independence axiom. In addition, the fact that the indifference curves are linear implies that an EU maximiser choosing A (respectively, B) from pair $\{A, B\}$, would also choose A (E) from pair $\{A, E\}$ and E (B) from pair $\{E, B\}$. This property is known as *betweenness*.

When applied to the lotteries within an M-M triangle, Becker *et al.*'s RP form of EUT implies that any pair of lotteries along any straight line of a certain gradient entails the same probability of choosing the safer option in the pair (S) over the riskier option (R), with that probability, $\Pr \text{Core}(S \succ R)$, reflecting the proportion of the DM's vNM functions favouring S .¹⁹ The DM's indifference curves are therefore represented by parallel straight lines with a slope such that $\Pr \text{Core}(S \succ R) = 0.5$. So, anyone choosing A from pair $\{A, B\}$ with some probability s , $0 \leq s \leq 1$, would also choose A from pair $\{A, E\}$, E from pair $\{E, B\}$ and C from pair $\{C, D\}$ with probability s .

Experimental work dating back to Allais (1953), Kahneman and Tversky (1979) and many others has shown that these predictions are often systematically violated. While many people tend to choose the safer option A in pairs such as $\{A, B\}$ and $\{A, E\}$, they tend to choose the riskier option D in pairs such as $\{C, D\}$. The reversal of modal preference between pairs $\{A, B\}$ and $\{C, D\}$ has come to be known as the *common ratio* (CR) effect, while the reversal of modal preference between $\{A, E\}$ and $\{C, D\}$ has come to be known as the *common consequence* (CC) effect. The betweenness property has also often been found to be violated, sometimes in the direction of the mixture being preferred over each of the two lotteries, sometimes in the opposite direction (e.g., Becker *et al.*, 1963b; Coombs and Huang, 1976; Bernasconi, 1994; see Blavatsky, 2006, for a recent overview).

We now look at BREUT's predictions in the M-M triangle. We start by applying the limiting case of unlimited sampling to analyse what happens when the probabilities of the best consequence are scaled down, as in the CR scenario (Theorem 1). We then derive more general properties about the shape of the indifference curves for cores of utility functions restricted to be either weakly concave or weakly convex (Theorem 2), which has implications for whether or not BREUT satisfies betweenness. By combining insights from Theorem 1 and Theorem 2, we also say more about the circumstances under which the CC effect can be obtained (Theorem 3). We conclude the section with a series of simulations illustrating the implications of BREUT when $d > 0$, which enable us to look into the behaviour of RT and $Conf^*$.

¹⁹ In each pair, the safer option is the one more to the south-west, the riskier the one more to the north-east. So, for example, A is the safer in $\{A, B\}$ and $\{A, E\}$, C is the safer in $\{C, D\}$. Throughout the rest of this section, we will use the convention that the first lottery in any pair $\{S, R\}$ is the safer and the second is the riskier.

Theorem 1

Let the core consist of $N > 1$ CRRA functions, i.e., functions of form $U_i(x) = x^{1-r_i}$, where $r_i < 1$ for all $i \in \{1, \dots, N\}$, and assume that $r_i \neq r_j$ for at least some $i, j \in \{1, \dots, N\}$. Take any lotteries $S = (x_m, p; 0, 1 - p)$ and $R = (x_h, q; 0, 1 - q)$ for which $x_h > x_m > 0$, $p \in (0, 1]$, $q \in (0, 1)$ and for which $\Pr(S \succ R) \rightarrow 0.5$ as $d \rightarrow 0$. Then, for any lottery $S' = (x_m, \sigma p; 0, 1 - \sigma p)$ and $R' = (x_h, \sigma q; 0, 1 - \sigma q)$ where $\sigma \in (0, 1)$, it must be that $\Pr(S' \succ R') \rightarrow 0$ as $d \rightarrow 0$.

Proof. See the Appendix.

Theorem 1 concerns lotteries of the form typically used in CR scenarios, $S = (x_m, p; 0, 1 - p)$ and $R = (x_h, q; 0, 1 - q)$, with $x_h > x_m > 0$. The commonly used case in which $p = 1$ is also allowed by Theorem 1. An example of lotteries of this form is presented in Figure 1, in which A, B, C and D fit the descriptions of S, R, S' and R' , respectively, with $p = 1, q = 0.8$ and $\sigma = 0.25$.

To illustrate the intuition, suppose that $\Pr(S \succ R) \rightarrow 0.5$ as $d \rightarrow 0$ (i.e., the decision maker is indifferent between S and R in the limit, so that $E[V(S, R)] = 0$) for some $p \leq 1$. The key result is that, for a core made of CRRA functions, scaling the probabilities of the best outcome of each lottery down by some factor σ always results in the DM having a strict preference for the scaled down risky lottery R' . Although Theorem 1 starts from perfect indifference between S and R in the scaled-up pair, an immediate corollary is that it will always be possible to obtain a strict preference in favour of S by slightly reducing q , the probability of winning x_h in R . That is, in the limit, a CR effect can always be obtained in which the DM has a strict preference for the safer lottery in the scaled-up pair and a strict preference for riskier lottery in the scaled-down pair. A preference reversal in the other direction does not occur for any lotteries of these forms.

Theorem 1 holds as long as there are at least two different CRRA utility functions in the core, without any further assumption about the core distribution, such as the degree of risk aversion implied by each function. The case of a continuous distribution (like the ones used in our simulations) can be approximated by taking an N that is sufficiently large.

Figure 2 illustrates the implications of Theorem 1. For convenience, pairs $\{S, R\}$ and $\{S', R'\}$ are drawn on (dashed) lines with the same gradient as the pairs of lotteries in Figure 1, which are also shown in the figure. S is drawn for some $p < 1$ (i.e., on the bottom edge of the triangle, but away from lottery A in the bottom-left corner). Because of the requirement that the DM is indifferent between S and R in the limit, S and R lie on the same indifference curve (the solid line connecting S and R). In the limiting case of BREUT, an indifference curve is the set of lotteries for which the core entails exactly the same mean CE . Theorem 1 implies that R' lies above the indifference curve passing through S' . In other words, indifference curves become flatter as one moves towards the bottom-right corner of the triangle, in line with the often discussed ‘fanning out’ pattern. Similarly, because the S in Figure 2 has been selected arbitrarily, the theorem implies that, to the left of S , indifference curves will be steeper. For a DM with indifference curves like those in Figure 2, $\Pr \text{Lim}(A, B) = 1$, $\Pr \text{Lim}(A, E) = 1$ and $\Pr \text{Lim}(C, D) = 0$, that is, she would display both the CR and the CC patterns with probability 1 in the limit.

[Insert Figure 2 here]

Figure 2 also serves to illustrate that, in the limiting case of BREUT, indifference curves are typically nonlinear. The nonlinearity implies that there is no guarantee that BREUT will satisfy betweenness in the limit. Theorem 2 sheds light on this aspect.

Theorem 2

Maintain the assumptions of Theorem 1 (i.e., the core consists of N CRRA functions, at least two of which are distinct). Take any lotteries $S = (x_m, p; 0, 1 - p)$ and $R = (x_h, q; 0, 1 - q)$ for which $x_h > x_m > 0$, $p \in (0, 1]$, $q \in (0, 1)$ and for which $\Pr(S \succ R) \rightarrow 0.5$ as $d \rightarrow 0$.

Suppose that all utility functions are (weakly) concave, i.e., $0 \leq r_i < 1$ for all $i \in \{1, \dots, N\}$. Then $\Pr(S \succ \omega S + (1 - \omega)R) \rightarrow 1$ and $\Pr(R \succ \omega S + (1 - \omega)R) \rightarrow 1$ as $d \rightarrow 0$, with $\omega \in (0, 1)$.²⁰

Suppose that all utility functions are (weakly) convex, i.e., $r_i \leq 0$ for all $i \in \{1, \dots, N\}$. Then $\Pr(S \succ \omega S + (1 - \omega)R) \rightarrow 0$ and $\Pr(R \succ \omega S + (1 - \omega)R) \rightarrow 0$ as $d \rightarrow 0$.

Proof. See the Appendix.

According to Theorem 2, if there are only (weakly) concave utility functions in the core, BREUT's indifference curves are always concave. This is the case depicted in Figure 2. Note that, since there must be at least two distinct utility functions in the core (as assumed in Theorem 1), the weak concavity requirement entails that there will be at least one strictly concave function. The key result is that any mixture of two lotteries, $S = (x_m, p; 0, 1 - p)$ and $R = (x_h, q; 0, 1 - q)$, that lie on the same indifference curve will be less preferred than either S or R . However, since the DM is assumed to be exactly indifferent between the two mixed lotteries, the degree of concavity of the indifference curves will typically be very small (unless the core is made of very extreme functions). This limits the room for observing violations of betweenness when sampling is limited (see our simulations below).²¹ If the core contains only (weakly) convex utility functions, then indifference curves will be convex and any mixture of S and R will be preferred to both. If there are both concave and convex utility functions in the core, the exact shape of the indifference curves will depend on the balance between concave and convex functions and on how extreme these are.²²

²⁰ Here and in the Appendix, we will use the notation $\omega S + (1 - \omega)R$ to indicate a lottery mixture of ω times the probability of the prizes inside the support of lottery S and $1 - \omega$ times the probability of the prizes in lottery R .

²¹ Note that Figure 2 is a sketch with an accentuated degree of concavity to illustrate the implications of our theorems more clearly.

²² Although Theorem 2 is stated in relation to pairs of lotteries one of which is located on the bottom edge and the other on the hypotenuse of the M-M triangle, this is only for simplicity of exposition. Theorem A.1 in the Appendix shows that an equivalent result holds for any pair of lotteries within the triangle.

An implication of Theorems 1 and 2 is that, while the CR effect will always be observed in the limit, for any core, the same is not guaranteed for the CC effect. But the effect will always be found in the limit if the core does not contain convex utility functions (the case illustrated in Figure 2), as detailed in Theorem 3.

Theorem 3

Maintain the assumptions of Theorem 1 (i.e., the core consists of N CRRA functions, at least two of which are distinct). Suppose that all utility functions are (weakly) concave, i.e., $0 \leq r_i < 1$ for all $i \in \{1, \dots, N\}$. Take any lotteries $T = (x_m, 1)$ and $Z = (x_h, q_1; x_m, q_2; 0, 1 - q_1 - q_2)$ for which $x_h > x_m > 0$, $q_1, q_2 \in (0, 1)$, $q_1 + q_2 < 1$ and $\Pr(T \succ Z) \rightarrow 0.5$ as $d \rightarrow 0$. Then it must be that for $T' = (x_m, 1 - q_2; 0, q_2)$ and $Z' = (x_h, q_1; 0, 1 - q_1)$, $\Pr(T' \succ Z') \rightarrow 0$ as $d \rightarrow 0$.

Proof. See the Appendix.

All these results hold for the limiting case of unlimited sampling, and we also know from Expression 6 that, for any $\{S, R\}$ pair, any pattern in $\Pr \text{Lim}(S, R)$ will eventually emerge in $\Pr(S, R)$ with sufficient sampling. The simulations reported in Table 7 explore the implications of the general form of BREUT for the four lottery pairs shown in Figure 1, using three values of d (0.1, 0.05 and 0.01), and setting $\alpha = 0.23$, $\beta = 1.0$. These levels of risk aversion, which are not atypical in experiments, have been chosen to obtain values of $\Pr \text{Core}(S \succ R)$ close to 0.5, in line with the assumptions of our theorems. This has the advantage of allowing us to illustrate the main implications of our model more clearly. Using positive values of d will also entail predictions for RT and Conf^* .

[Insert Table 7 here]

We start by looking at the $E[V(S, R)]$ column of Table 7, which links up to the limiting case. For all pairs except $\{C, D\}$, $E[V(S, R)]$ is positive, so we can expect that in each of these pairs the safer option will be chosen with probability 1 in the limit. The negative value for $\{C, D\}$ implies that in the limit the riskier option will be chosen with probability 1, in line with both the CR and CC effects. Because of the signs of $E[V(S, R)]$ and the fact that $\Pr \text{Core}(S \succ R) = 0.51$, Expression 6 implies that, with limited sampling, there will be a reversal of modal choice in both the CR and CC scenarios, as we see in all panels of Table 7. In addition, $\Pr(S \succ R)$ moves further away from 0.5 as d decreases and the DM samples more. As can be expected, as d gets smaller, RT s increase on average; correspondingly, the DM makes her choices with higher confidence.

In all the cases presented in Table 7, there are no violations of probabilistic betweenness (in all three panels, the choice probabilities for pairs $\{A, B\}$, $\{A, E\}$ and $\{E, B\}$ are virtually identical). From the $E[V(S, R)]$ column, we know that $\Pr \text{Lim}(A, B) = 1$, $\Pr \text{Lim}(A, E) = 1$ and $\Pr \text{Lim}(E, B) = 1$. We know from Theorem 2 that in order for violations of betweenness to be guaranteed in the limit, the core of BREUT must consist of either all (weakly) concave or all (weakly) convex utility functions and $E[V(S, R)]$ must be (close to) zero.

Neither of these conditions is satisfied in the simulations presented in Table 7 (the core contains a mix of concave and convex functions, and $E[V(A, B)]$ is the largest of the three mean CE differences). Because the curvature of the indifference curves is typically not very pronounced, even when these conditions are satisfied it takes a lot of sampling (i.e., very low values of d) for BREUT to produce systematic violations of betweenness of any significance. Table 8 illustrates.

[Insert Table 8 here]

The two panels of Table 8 consider two stripped down versions of the model, in which there are just two utility functions in the core, which are sampled with equal probability. The functions in the top panel are weakly concave ($r = 0$ and $r \approx 0.72$), while those in the second panel are weakly convex ($r = 0$ and $r \approx -3.35$). In each case, the two values of r are chosen so that the DM is indifferent in the limit, i.e. $E[V(S, R)] = 0$. There are just two pairs, the equivalents of $\{A, B\}$ and $\{E, B\}$ in Figure 1, but A is adjusted to reach indifference in the limit, so we use $A' = (25, 1)$ in the top panel and $A'' = (35, 1)$ in the bottom panel. We only look at two pairs because the first pair is always such that both $\Pr \text{Core}(S \succ R) = 0.50$ and $\Pr(S \succ R) = 0.50$. Therefore, violations of betweenness can be detected if E is chosen over either A' (A'') or B with a probability different from 0.5. In line with Theorem 2, we see that $\Pr(E \succ B) = 0.48$ when the functions are weakly concave: that is, E is less preferred than both A' and B more than fifty percent of the time. When the functions are weakly convex, $\Pr(E \succ B) = 0.53$: that is, E is preferred to both A'' and B more than fifty percent of the time. In spite of d being as low as 0.01, these effects are rather small and require substantial amounts of sampling to become evident. Because $\Pr(S \succ R)$ is close to 0.5 in all cases and sampling is substantial, RTs tend to be high and $Conf^*$ values low, reaching similar levels for weakly concave and weakly convex functions. Overall, although BREUT does not satisfy the betweenness property, the implied violations are likely to be small and not easily detectable even with considerable amounts of sampling.

4. Discussion and Conclusions

In this final section, we consider where BREUT stands in relation to various other models, we discuss some of its limitations and identify possible extensions, we expand on some broader implications of our findings and we offer some concluding remarks.

4.1. BREUT vs. Other Stochastic Models of Risky Choice

BREUT is, in effect, a sequential-sampling formulation of the RP specification in Becker *et al.*'s (1963a) random utility model for wagers. While Becker *et al.* modelled choice as if a DM acts on the basis of drawing a single $u(\cdot)$ from a distribution, BREUT supposes that deliberation involves a series of such draws. When the single-draw RP model is applied to CR and CC pairs, it entails that the probability of choosing the safe option is the same in both pairs. However, in BREUT, sequential sampling allows the probability of safe choices to vary, sometimes to the extent of reversing the modal choice.

These reversals also set BREUT apart from what is arguably the most standard way of incorporating stochasticity into choice under risk: the Fechnerian version of EUT (e.g., Fechner, 1860/1966; Hey and Orme, 1994). In Fechnerian models, a DM is assumed to have a single vNM utility function which gives a ‘true’ utility difference for one option over another, but her choices are stochastic because some extraneous noise or error component is added. Sometimes the error will reinforce her true preference, but sometimes it may work in the opposite direction and – if sufficiently large – may outweigh the true difference so that on some occasions the DM chooses the truly-less-preferred option. Moreover, the frequency with which this occurs may vary from pair to pair, depending on the relationship between the true utility difference and the variance of the error term.

So, an EU maximiser whose $u(\cdot)$ favours the safer A in the scaled-up CR pair $\{A, B\}$ will also truly prefer the safer C in the scaled-down $\{C, D\}$ pair; but if the utility difference in the $\{C, D\}$ pair is smaller than in the $\{A, B\}$ pair, homoscedastic Fechnerian noise is liable to overturn that difference more often, resulting in a lower probability of choosing C than of choosing A . This pattern is in the direction of the usual CR effect. However, if – as is usually assumed – the median of the error term is zero, the Fechner model with a single $u(\cdot)$ at the core cannot produce a reversal of the modal choice between the two pairs: the most that can happen is that as the EU difference becomes small relative to the variance of the error term, the probability of choosing the truly-more-preferred option tends to decrease towards, but never falls below, 0.5. In addition, since the EU differences for pairs such as $\{A, E\}$ and $\{C, D\}$ are of similar magnitude, the model in its simple standard form has no way of producing the CC effect.

In addition, the Fechner model would in some cases lead us to expect many more violations of FOSD than are typically observed when dominance is transparent and easy to detect (see Bardsley *et al.*, 2010, chapter 7, for a discussion). In all these respects, the standard Fechner model is different from, and empirically less well supported than, BREUT. Some refinements of the standard Fechnerian approach have tried to fix some of these problems, for instance by introducing specific assumptions about the distribution of noise (e.g., Blavatsky, 2007, 2011; Wilcox, 2011). Generally, however, even these more elaborate Fechnerian error stories offer no account for the differences in response times and in confidence which, as we have seen, are intrinsic to a deliberative process such as the one modelled in BREUT and are also consistent with empirical evidence. So, while we do not deny that extraneous noise may have some effect upon the choices we observe, our analysis does not depend on it and we have derived all of our results as if there were no such additional component.²³

The model that bears the closest resemblance to BREUT in the decision-making literature is one of the intermediate stages considered by Busemeyer and Townsend (1993) in their derivation of Decision Field Theory (DFT). In that paper, DFT is presented as the result of seven consecutive stages that evolve from deterministic Subjective EUT (Stage 1) to the fully-fledged version of the DFT model (Stage 7). Stage 2

²³ Loomes (2014) argues that extraneous noise may sometimes have a substantial additional effect on the data generated by choice experiments, and that the coexistence of such noise with intrinsic variability of the kind we have focused on may complicate and compromise the testing of core theories. However, in this paper we set such complications to one side.

modifies standard Subjective EUT by introducing fluctuating attention weights. Stage 3 embeds this model in a sequential-sampling framework. At first glance it might seem that Stage 3 DFT is much the same as the model we are proposing, but there are four key differences.

First, in Stage 3 DFT, the specification assumes that the stochastic variability is produced by fluctuations in the probabilities of comparing different pairs of payoffs, rather than fluctuations in the subjective values of those payoffs. Second, the probabilities are transformed in a way that deviates from EUT. This alone is sufficient to produce non-EU patterns of choice and makes the model unsuitable for exploring the consequences of deliberation based on EU preferences (which is one of the main goals of BREUT). Third, the magnitudes being accumulated in Stage 3 DFT are not differences in *CEs*, but differences between the utilities of the alternative options after they have been weighted by transformed probabilities. Fourth, the sequential-sampling process in DFT takes the form of a Markov process with absorbing thresholds, whereas BREUT uses a stopping rule based on a sequential statistical test. This allows BREUT to make predictions about confidence, which standard DFT (either Stage 3 or the fully-fledged version) is silent about.

After Stage 3, there are various additional stages of elaboration, so that the full (Stage 7) DFT model is considerably more complex than BREUT and involves a larger number of parameters. As a result, this full form of DFT can explain some aspects of decision behaviour that BREUT cannot, such as violations of FOSD for gambles with negatively correlated consequences (see Busemeyer and Townsend, 1993, p. 447).

4.2. Possible Extensions

Our objective in this paper is not to try to explain all of the regularities that have been documented in decades of research in decision making under risk or try to reconcile them with an EUT core, but rather to explore the implications of embedding standard EUT in a boundedly rational deliberation process. As we indicated, our general modelling approach is not specific to EUT but can be extended to any other core theory that is able to generate *CE* differences between two options. Tracing the implications of embedding various non-EU core theories in a sequential-sampling framework will require a substantial programme of research that goes beyond the scope of this paper, but it is a promising line of research for the future. Incorporating a non-EU core would expand the range of regularities that can be explained using accumulation and stopping mechanisms like the one adopted in BREUT.

For example, one way of accommodating more pronounced betweenness violations would be to include some form of probability weighting in the model. In addition, while our approach focused on the domain of gains, it would in principle be possible to also study behaviour in the domain of losses or for mixed lotteries. This could be achieved by incorporating a reference point and allowing for asymmetric attitudes to gains and losses. Such extensions would require sampling over more than one dimension – e.g., risk aversion and degree of probability weighting, or risk aversion and loss aversion – and modelling the joint distribution of these parameters.

Another direction in which the modelling of deliberation could usefully be developed concerns the process by which individuals generate *matching* or *equivalence* responses. In this context, a task widely used

in experiments and surveys is the request to respondents to provide a best estimate of their willingness-to-pay (WTP) or willingness-to-accept (WTA) valuations.

In some studies, these are elicited using *multiple price lists*, which may be regarded as a sequence of pairwise choices. However, asking DMs to work through an ordered list is likely to entail that the choices within the list are not treated independently. Extending sequential sampling models to such tasks would therefore require us to address how non-independence should be modelled, and what the possible implications might be.

In other studies, the elicitation of WTP or WTA is done via *open-ended* questions, which require the participant to simply provide a single monetary value as a response. One way of adapting our model to produce responses to this latter procedure might involve sampling utility functions and accumulating the resulting *CEs*, until an appropriate stopping rule prompts the individual to state a valuation based on the information acquired up to that point. However, there are other ways of modelling the process behind equivalence judgments,²⁴ the exploration of which is a challenge that we leave for future research.

4.3 Broader Implications and Conclusions

We conclude by looking at some broader issues raised by our results.

Kahneman *et al.* (2016) have recently stressed that variability in decision making can be extremely costly to organisations. They provide examples of how key DMs, faced with the same scenarios, are liable to arrive at different judgments on different occasions, making it difficult to ensure reliability in crucial areas of organisational activity. They also note that the noisy nature of organizational processes is often not recognized by managers, who tend to assume that a particular procedure will consistently lead to the same outcome. This raises the question of how variability could be reduced to improve consistency, and Kahneman *et al.* make a number of proposals to achieve that end (e.g., ‘noise audits’, provision of checklists and algorithms). In addition, our model suggests that it would be useful to elicit decision makers’ assessments of the confidence they have in decisions, with a view to identifying those where it is more probable that different options may be chosen. It may also help (when possible without the DMs being aware) to make the same individuals repeat their decisions, so that variability, and its associated costs, can be measured at the individual level.

A distinctive implication of our results is that the evidence accumulation process affects the extent to which patterns of choice can be said to reveal directly the structure of underlying preferences.²⁵ In BREUT, the DM’s core preferences consist exclusively of vNM utility functions, but deliberation produces patterns of final choices that systematically depart from those implied by independence, betweenness and strong stochastic transitivity. When modelling the distribution of those underlying preferences, we used a multiple-selves analogy such that an individual’s eventual decision was depicted as the aggregate of a

²⁴ Within the framework of DFT, Johnson and Busemeyer (2005, pp. 844-6) outlined a ‘sequential value-matching’ algorithm that represented valuation as a somewhat indirect process, a form of iterative binary choice between the option being valued and a series of possible sure amounts.

²⁵ For a related result in the context of inferring preferences from monetary certainty equivalents, see Wilcox (2016).

sample of those inner selves' strengths of preference expressed in certainty equivalent form. There is an analogy here with cases in welfare economics where public policy is based on aggregating values over a population of different individuals with heterogeneous preferences. For example, suppose one were interested in valuing different degrees of risk reduction provided by two alternative interventions using measures of WTP averaged over a sample of individuals. The between-person analogue of Theorem 1 suggests that, even if every person in the sample were a deterministic EU maximiser, there could be a reversal of the aggregate preference over the two interventions simply based on whether the probabilities involved were scaled up or down, even though none of the individual preferences in the population implied such a reversal.

Our results also have implications for experimental investigations of decision making. Recognising the probabilistic nature of choice, one experimental design strategy has involved trying to collect repeated responses to the same decision problems to estimate choice variability. This is a methodology that can lead to complications (e.g., fatigue, boredom, remembering previous decisions, etc.), especially if the number of repetitions is high. However, because response times and confidence are systematically related to choice probabilities, investigators may be able to collect measures of those variables and use them in conjunction with fewer repetitions to produce data sets that are less vulnerable to such complications. That implication is not specific to an EUT core and *RTs* (which are easy to collect in computerised experiments) may therefore be a useful adjunct to many studies testing competing core theories.

To conclude, the BREUT model set out in this paper can be seen as an illustration of the importance of taking the deliberation process into account when modelling decisions, as advocated by Simon many years ago. Allocating theoretical and empirical effort towards exploring the 'production of decisions' is, we suggest, likely to open fruitful avenues for future research.

Appendix

The proofs below make use of the following observation. For any two lotteries A and B , if $\Pr(A > B) \rightarrow 1$ as $d \rightarrow 0$, then $E[V(A, B)] > 0$, and so the expected CE of A is greater than the expected CE of B (i.e., $E(CE_A) > E(CE_B)$), where, for any lottery L , $E(CE_L)$ denotes its expected CE . If $\Pr(A > B) \rightarrow 0.5$ as $d \rightarrow 0$, then $E[V(A, B)] = 0$ and $E(CE_A) = E(CE_B)$. Recall also that the mixture operator is defined in the usual way, i.e., if $A(x)$ and $B(x)$ denote the probability of reaching prize x in lotteries A and B , respectively, then $(\omega A + (1 - \omega)B)(x) = \omega A(x) + (1 - \omega)B(x)$ for all x .

Theorem 1: See main text.

Proof. For convenience, define $\theta_i = \frac{1}{1-r_i}$, so that $U_i(x) = x^{1/\theta_i}$, with $\theta_i > 0$. Without loss of generality, assume the θ_i 's are ordered according to $\theta_1 > \dots > \theta_N$. (Note that, while in principle it may be that $\theta_i = \theta_j$ for some distinct i, j , there are at least one i and one j such that $\theta_i \neq \theta_j$. So, the above ordering can be done by replacing N with the cardinality of the strict ordering, which is of at least 2.) Let f_i denote the probability of θ_i occurring. We proceed by contradiction. Suppose that the result does not hold, so that $\Pr(S > R) \rightarrow 0.5$ and $\Pr(S' > R')$ does not tend to 0 as $d \rightarrow 0$. First, $E(CE_S) = E(CE_R)$ and hence

$$x_m \sum_{i=1}^N f_i p^{\theta_i} = x_h \sum_{i=1}^N f_i q^{\theta_i}, \text{ which implies } x_m = \frac{x_h \sum_{i=1}^N f_i q^{\theta_i}}{\sum_{i=1}^N f_i p^{\theta_i}}.$$

Second, since it must be that $\Pr(S' > R')$ tends to a number greater than 0 as $d \rightarrow 0$ (and specifically 0.5 or 1), then $E(CE_{S'}) \geq E(CE_{R'})$, and hence

$$x_m \sum_{i=1}^N f_i (\sigma p)^{\theta_i} \geq x_h \sum_{i=1}^N f_i (\sigma q)^{\theta_i}, \text{ which implies } x_m \geq \frac{x_h \sum_{i=1}^N f_i (\sigma q)^{\theta_i}}{\sum_{i=1}^N f_i (\sigma p)^{\theta_i}}.$$

Combining the two, we obtain:

$$\frac{\sum_{i=1}^N f_i q^{\theta_i}}{\sum_{i=1}^N f_i p^{\theta_i}} \geq \frac{\sum_{i=1}^N f_i (\sigma q)^{\theta_i}}{\sum_{i=1}^N f_i (\sigma p)^{\theta_i}},$$

where we have divided by x_h on both sides. Noting that $q < p$, we write $q = bp$, where $b = q/p \in (0, 1)$.

We therefore have:

$$\begin{aligned} & \frac{\sum_{i=1}^N f_i (bp)^{\theta_i}}{\sum_{i=1}^N f_i p^{\theta_i}} \geq \frac{\sum_{i=1}^N f_i (\sigma bp)^{\theta_i}}{\sum_{i=1}^N f_i (\sigma p)^{\theta_i}}, \\ \Rightarrow & \sum_{j=1}^N f_j (bp)^{\theta_j} \sum_{i=1}^N f_i (\sigma p)^{\theta_i} \geq \sum_{j=1}^N f_j (\sigma bp)^{\theta_j} \sum_{i=1}^N f_i p^{\theta_i}. \end{aligned}$$

So:

$$(f_1 b^{\theta_1} p^{\theta_1} + \dots + f_N b^{\theta_N} p^{\theta_N})(f_1 \sigma^{\theta_1} p^{\theta_1} + \dots + f_N \sigma^{\theta_N} p^{\theta_N}) \geq$$

$$(f_1 \sigma^{\theta_1} b^{\theta_1} p^{\theta_1} + \dots + f_N \sigma^{\theta_N} b^{\theta_N} p^{\theta_N})(f_1 p^{\theta_1} + \dots + f_N p^{\theta_N}).$$

Factoring out the left-hand side (LHS) and the right-hand side (RHS) of the equation:

$$\begin{aligned} & f_1^2 b^{\theta_1} \sigma^{\theta_1} p^{2\theta_1} + f_1 f_2 b^{\theta_1} \sigma^{\theta_2} p^{\theta_1 + \theta_2} + \dots + f_1 f_N b^{\theta_1} \sigma^{\theta_N} p^{\theta_1 + \theta_N} + \\ & f_2 f_1 b^{\theta_2} \sigma^{\theta_1} p^{\theta_1 + \theta_2} + f_2^2 b^{\theta_2} \sigma^{\theta_2} p^{2\theta_2} + \dots + f_2 f_N b^{\theta_2} \sigma^{\theta_N} p^{\theta_2 + \theta_N} + \\ & \quad + \dots + \\ & f_N f_1 b^{\theta_N} \sigma^{\theta_1} p^{\theta_1 + \theta_N} + f_N f_2 b^{\theta_N} \sigma^{\theta_2} p^{\theta_2 + \theta_N} + \dots + f_N^2 b^{\theta_N} \sigma^{\theta_N} p^{2\theta_N} \\ & \geq \\ & f_1^2 b^{\theta_1} \sigma^{\theta_1} p^{2\theta_1} + f_1 f_2 b^{\theta_1} \sigma^{\theta_1} p^{\theta_1 + \theta_2} + \dots + f_1 f_N b^{\theta_1} \sigma^{\theta_1} p^{\theta_1 + \theta_N} + \\ & f_2 f_1 b^{\theta_2} \sigma^{\theta_2} p^{\theta_1 + \theta_2} + f_2^2 b^{\theta_2} \sigma^{\theta_2} p^{2\theta_2} + \dots + f_2 f_N b^{\theta_2} \sigma^{\theta_2} p^{\theta_2 + \theta_N} + \\ & \quad + \dots + \\ & f_N f_1 b^{\theta_N} \sigma^{\theta_N} p^{\theta_1 + \theta_N} + f_N f_2 b^{\theta_N} \sigma^{\theta_N} p^{\theta_2 + \theta_N} + \dots + f_N^2 b^{\theta_N} \sigma^{\theta_N} p^{2\theta_N}. \end{aligned}$$

Canceling out the common terms on the LHS and the RHS side and grouping the terms that include $f_i f_j$ with those that include $f_j f_i$, we obtain:

$$\begin{aligned} & f_1 f_2 p^{\theta_1 + \theta_2} (b^{\theta_1} \sigma^{\theta_2} + b^{\theta_2} \sigma^{\theta_1}) + f_1 f_3 p^{\theta_1 + \theta_3} (b^{\theta_1} \sigma^{\theta_3} + b^{\theta_3} \sigma^{\theta_1}) + \dots + f_1 f_N p^{\theta_1 + \theta_N} (b^{\theta_1} \sigma^{\theta_N} + b^{\theta_N} \sigma^{\theta_1}) + \\ & \quad f_2 f_3 p^{\theta_2 + \theta_3} (b^{\theta_2} \sigma^{\theta_3} + b^{\theta_3} \sigma^{\theta_2}) + \dots + \\ & f_{N-1} f_N p^{\theta_{N-1} + \theta_N} (b^{\theta_{N-1}} \sigma^{\theta_N} + b^{\theta_N} \sigma^{\theta_{N-1}}) \geq f_1 f_2 p^{\theta_1 + \theta_2} (b^{\theta_1} \sigma^{\theta_1} + b^{\theta_2} \sigma^{\theta_2}) + f_1 f_3 p^{\theta_1 + \theta_3} (b^{\theta_1} \sigma^{\theta_1} + \\ & \quad b^{\theta_3} \sigma^{\theta_3}) + f_2 f_3 p^{\theta_2 + \theta_3} (b^{\theta_2} \sigma^{\theta_2} + b^{\theta_3} \sigma^{\theta_3}) + \dots + f_{N-1} f_N p^{\theta_{N-1} + \theta_N} (b^{\theta_{N-1}} \sigma^{\theta_{N-1}} + b^{\theta_N} \sigma^{\theta_N}), \end{aligned}$$

which can be written as:

$$\sum_{j=i+1}^N \sum_{i=1}^{N-1} f_i f_j p^{\theta_i + \theta_j} (b^{\theta_i} \sigma^{\theta_j} + b^{\theta_j} \sigma^{\theta_i}) \geq \sum_{j=i+1}^N \sum_{i=1}^{N-1} f_i f_j p^{\theta_i + \theta_j} (b^{\theta_i} \sigma^{\theta_i} + b^{\theta_j} \sigma^{\theta_j}).$$

But notice that for any $j > i$, $f_i f_j p^{\theta_i + \theta_j} (b^{\theta_i} \sigma^{\theta_j} + b^{\theta_j} \sigma^{\theta_i}) < f_i f_j p^{\theta_i + \theta_j} (b^{\theta_i} \sigma^{\theta_i} + b^{\theta_j} \sigma^{\theta_j})$, as it is implied by $b^{\theta_i} \sigma^{\theta_j} + b^{\theta_j} \sigma^{\theta_i} < b^{\theta_i} \sigma^{\theta_i} + b^{\theta_j} \sigma^{\theta_j}$, which is implied by $b^{\theta_i} (\sigma^{\theta_j} - \sigma^{\theta_i}) < b^{\theta_j} (\sigma^{\theta_j} - \sigma^{\theta_i})$, which is implied by $\sigma^{\theta_j} - \sigma^{\theta_i} > 0$ and by $b^{\theta_i} < b^{\theta_j}$, which are both true because $\sigma, b \in (0,1)$ and $\theta_i > \theta_j > 0$. Since this is true for each of the comparisons of the matching terms in the LHS and the RHS sums, it follows that:

$$\sum_{j=i+1}^N \sum_{i=1}^{N-1} f_i f_j p^{\theta_i + \theta_j} (b^{\theta_i} \sigma^{\theta_j} + b^{\theta_j} \sigma^{\theta_i}) < \sum_{j=i+1}^N \sum_{i=1}^{N-1} f_i f_j p^{\theta_i + \theta_j} (b^{\theta_i} \sigma^{\theta_i} + b^{\theta_j} \sigma^{\theta_j}),$$

which is a contradiction and completes the proof. ■

Below we provide a more general version of Theorem 2 than the one included in the main text. By proving the more general theorem we also provide a proof for Theorem 2.

Theorem A.1 (generalized Theorem 2)

Maintain the assumptions of Theorem 1 (i.e., the core consists of N CRRA functions, at least two of which are distinct). Take any distinct lotteries $S = (x_h, p_1; x_m, p_2; 0, 1 - p_1 - p_2)$ and $R = (x_h, q_1; x_m, q_2; 0, 1 - q_1 - q_2)$ for which $x_h > x_m > 0, p_1, p_2, q_1, q_2 \in [0, 1], p_1 + p_2 < 1, q_1 + q_2 < 1$ and for which $\Pr(S \succ R) \rightarrow 0.5$ as $d \rightarrow 0$. Suppose that all utility functions are (weakly) concave, i.e., $0 \leq r_i < 1$ for all $i \in \{1, \dots, N\}$. Then $\Pr(S \succ \omega S + (1 - \omega)R) \rightarrow 1$ and $\Pr(R \succ \omega S + (1 - \omega)R) \rightarrow 1$ as $d \rightarrow 0$, where $\omega \in (0, 1)$.

Suppose that all utility functions are (weakly) convex, i.e., $r_i \leq 0$ for all $i \in \{1, \dots, N\}$. Then $\Pr(S \succ \omega S + (1 - \omega)R) \rightarrow 0$ and $\Pr(R \succ \omega S + (1 - \omega)R) \rightarrow 0$ as $d \rightarrow 0$.

Proof. We only prove the case for $0 \leq r_i < 1$ for all $i \in \{1, \dots, N\}$; an analogous proof holds for $r_i \leq 0$ for all $i \in \{1, \dots, N\}$. Adopt the same notation as in the proof of Theorem 1, with the added restriction here that $\theta_i \geq 1$ for all i . Moreover, since there are at least two distinct θ 's in the support, it must be that $\theta_i > 1$ for at least one i , i.e., it must be that at least one utility function is strictly concave.

Since $\Pr(S \succ R) \rightarrow 0.5$ as $d \rightarrow 0$, it must be that $E(CE_S) = E(CE_R)$ and hence that

$$\sum_{i=1}^N f_i(p_1 x_h^{1/\theta_i} + p_2 x_m^{1/\theta_i})^{\theta_i} = \sum_{i=1}^N f_i(q_1 x_h^{1/\theta_i} + q_2 x_m^{1/\theta_i})^{\theta_i}.$$

Moreover, since $\omega S + (1 - \omega)R = (x_h, \omega p_1 + (1 - \omega)q_1; x_m, \omega p_2 + (1 - \omega)q_2; 0, 1 - \omega(p_1 + p_2) - (1 - \omega)(q_1 + q_2))$,

$$\begin{aligned} E(CE_{\omega S + (1 - \omega)R}) &= \sum_{i=1}^N f_i((\omega p_1 + (1 - \omega)q_1)x_h^{1/\theta_i} + (\omega p_2 + (1 - \omega)q_2)x_m^{1/\theta_i})^{\theta_i} \\ &= \sum_{i=1}^N f_i(\omega(p_1 x_h^{1/\theta_i} + p_2 x_m^{1/\theta_i}) + (1 - \omega)(q_1 x_h^{1/\theta_i} + q_2 x_m^{1/\theta_i}))^{\theta_i}. \end{aligned}$$

Since $\theta_i \geq 1$ for all i and $\theta_i > 1$ for at least one i , it follows from Jensen's Inequality for weakly and strictly convex functions that:

$$\begin{aligned} &(\omega(p_1 x_h^{1/\theta_i} + p_2 x_m^{1/\theta_i}) + (1 - \omega)(q_1 x_h^{1/\theta_i} + q_2 x_m^{1/\theta_i}))^{\theta_i} \\ &\leq \omega(p_1 x_h^{1/\theta_i} + p_2 x_m^{1/\theta_i})^{\theta_i} + (1 - \omega)(q_1 x_h^{1/\theta_i} + q_2 x_m^{1/\theta_i})^{\theta_i} \end{aligned}$$

for all i , with strict inequality for at least one i for which $\theta_i > 1$. The reason the inequality is strict for at

least one θ_i is that equality would only hold if for all $\theta_i, p_1 x_h^{1/\theta_i} + p_2 x_m^{1/\theta_i} = q_1 x_h^{1/\theta_i} +$

$q_2 x_m^{1/\theta_i}$, i.e., $(x_h/x_m)^{1/\theta_i}(p_1 - q_1) = q_2 - p_2$. But this is impossible because there are at least two

distinct θ 's (i.e., $\theta_j \neq \theta_k$ for at least some $j, k \in \{1, \dots, N\}$), and so $(x_h/x_m)^{1/\theta_i}$, and hence the LHS, is

different for the distinct $\theta_j \neq \theta_k$, while the RHS maintains the same value, and so equality can hold at most for one θ (and note that in the case where $N = 2$, $\theta_1 > 1$, and $\theta_2 = 1$, then clearly from $E(CE_S) = E(CE_R)$, $p_1x_h^{1/\theta_1} + p_2x_m^{1/\theta_1} = q_1x_h^{1/\theta_1} + q_2x_m^{1/\theta_1}$ could not hold for either θ_1 or θ_2 . Hence, even in that case, strict inequality would still hold for $\theta_1 > 1$). Therefore, combining all terms, we obtain:

$$\begin{aligned}
E(CE_{\omega S + (1-\omega)R}) &= \sum_{i=1}^N f_i(\omega(p_1x_h^{1/\theta_i} + p_2x_m^{1/\theta_i}) + (1-\omega)(q_1x_h^{1/\theta_i} + q_2x_m^{1/\theta_i}))^{\theta_i} \\
&< \sum_{i=1}^N f_i(\omega(p_1x_h^{1/\theta_i} + p_2x_m^{1/\theta_i})^{\theta_i} + (1-\omega)(q_1x_h^{1/\theta_i} + q_2x_m^{1/\theta_i})^{\theta_i}) \\
&= \omega \sum_{i=1}^N f_i(p_1x_h^{1/\theta_i} + p_2x_m^{1/\theta_i})^{\theta_i} + (1-\omega) \sum_{i=1}^N f_i(q_1x_h^{1/\theta_i} + q_2x_m^{1/\theta_i})^{\theta_i} \\
&= \omega E(CE_S) + (1-\omega)E(CE_R) = E(CE_S) = E(CE_R),
\end{aligned}$$

where we have used that $E(CE_S) = \sum_{i=1}^N f_i(p_1x_h^{1/\theta_i} + p_2x_m^{1/\theta_i})^{\theta_i}$ and $E(CE_R) = \sum_{i=1}^N f_i(q_1x_h^{1/\theta_i} + q_2x_m^{1/\theta_i})^{\theta_i}$. We have therefore shown that $E(CE_{\omega S + (1-\omega)R}) < E(CE_S) = E(CE_R)$, from which it follows that $\Pr(S > \omega S + (1-\omega)R) \rightarrow 1$ and $\Pr(R > \omega S + (1-\omega)R) \rightarrow 1$ as $d \rightarrow 0$. This concludes the proof. The proof for $r_i \leq 0$ is analogous, but uses Jensen's Inequality for concave and strictly concave functions instead. ■

Theorem 3: See main text.

Proof. Adopt the same notation as in the proofs of Theorems 1 and A.1. We make use of both theorems to prove the result. First, it follows from $\Pr(T > Z) \rightarrow 0.5$ as $d \rightarrow 0$ that $E(CE_T) = E(CE_Z)$. Consider now the lottery $G = (x_h, \frac{q_1}{1-q_2}; 0, \frac{1-q_1-q_2}{1-q_2})$, which is at the intersection between the hypotenuse and the line segment connecting lotteries T and Z in the corresponding Marschak-Machina diagram. Note that $Z = q_2T + (1 - q_2)G$, since $q_2T + (1 - q_2)G = q_2(x_m, 1) + (1 - q_2)(x_h, \frac{q_1}{1-q_2}; 0, \frac{1-q_1-q_2}{1-q_2}) = (x_h, q_1; x_m, q_2; 0, 1 - q_1 - q_2)$. By Theorem A.1, it cannot be that $E(CE_G) = E(CE_T)$, as this would imply that $E(CE_Z) < E(CE_T)$. It is also obvious from monotonicity that $E(CE_G)$ cannot be strictly less than $E(CE_T)$, which would also imply that $E(CE_Z) < E(CE_T)$. It must therefore be that $E(CE_G) > E(CE_T)$. Now from Theorem 1, we know that if $E(CE_T) = E(CE_G)$ holds (i.e., $\Pr(T > G) \rightarrow 0.5$ as $d \rightarrow 0$), then $E(CE_T) < E(CE_Z)$ (i.e., $\Pr(T' > Z') \rightarrow 0$ as $d \rightarrow 0$). This can be seen by taking p in Theorem 1 to be 1, q to be $\frac{q_1}{1-q_2}$, and σ to be $1 - q_2$, and replacing S with T , R with G , S' with T' and R' with Z' . It is then clear by monotonicity that $E(CE_G) > E(CE_T)$ also implies that $\Pr(T' > Z') \rightarrow 0$ as $d \rightarrow 0$. ■

References

- Acemoglu, D. and Scott, A. (1994). Consumer confidence and rational expectations: Are agents' beliefs consistent with the theory? *Economic Journal*, 104, 1–19.
- Achtziger, A. and Alós-Ferrer, C. (2014). Fast or rational? A response-time study of Bayesian updating. *Management Science*, 60(4), 923–938.
- Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: Critique des postulats et axiomes de l'école Américaine. *Econometrica*, 21, 503–546.
- Alós-Ferrer, C. and Strack, F. (2014). From dual processes to multiple selves: Implications for economic behavior. *Journal of Economic Psychology*, 41, 1–11.
- Bardsley, N., Cubitt, R., Loomes, G., Moffatt, P., Starmer, C. and Sugden, R. (2010). *Experimental Economics: Rethinking the Rules*, Princeton University Press.
- Barsky, R. B. and Sims, E. R. (2012). Information, animal spirits, and the meaning of innovations in consumer choice. *American Economic Review*, 102(4), 1343–1377.
- Baucells, M. and Villasís, A. (2010). Stability of risky preferences and the reflection effect of prospect theory. *Theory and Decision*, 68, 193–211.
- Becker, G. M., DeGroot, M. H. and Marschak, J. (1963a). Stochastic models of choice behavior. *Behavioral Science*, 8, 41–55.
- Becker, G. M., DeGroot, M. H. and Marschak, J. (1963b). An experimental study of some stochastic models of wagers. *Behavioral Science*, 8, 199–202.
- Bernasconi, M. (1994). Nonlinear preferences and two-stage lotteries: Theories and evidence. *Economic Journal*, 104, 54–70.
- Binmore, K. (2009). Interpersonal comparison of utility. In D. Ross and H. Kincaid (Eds.), *The Oxford Handbook of Philosophy of Economics* (pp. 540–559). Oxford University Press: New York.
- Birnbaum, M. H. (2008). New paradoxes of risky decision making. *Psychological Review*, 115, 463–501.
- Birnbaum, M. H. and Jou, J. (1990). A theory of comparative response times and 'difference' judgments. *Cognitive Psychology*, 22, 184–210.
- Blavatskyy, P. (2006). Violations of betweenness or random errors? *Economics Letters*, 91, 34–38.
- Blavatskyy, P. (2007). Stochastic expected utility theory. *Journal of Risk and Uncertainty*, 34 (3), 259–286.
- Blavatskyy, P. R. (2011). A model of probabilistic choice satisfying first-order stochastic dominance. *Management Science*, 57, 542–548.
- Blavatskyy, P. R., and Pogrebna, G. (2010). Models of stochastic choice and decision theories: Why both are important for analyzing decisions. *Journal of Applied Econometrics*, 25, 963–986.
- Brandstätter, E., Gigerenzer, G. and Hertwig, R. (2006). The priority heuristic: Making choices without trade-offs. *Psychological Review*, 113, 409–432.
- Busemeyer, J. R. and Townsend, J. T. (1993). Decision field theory: A dynamic-cognitive approach to decision making in an uncertain environment. *Psychological Review*, 100, 432–459.
- Butler, D. and Loomes, G. (1988). Decision difficulty and imprecise preferences. *Acta Psychologica*, 68, 183–196.
- Butler, D., Isoni, A., Loomes, G. and Tsutsui, K. (2014a). Beyond choice: Investigating the sensitivity and validity of strength of preference measures. *Experimental Economics*, 17, 537–563.
- Butler, D., Isoni, A., Loomes, G. and Navarro-Martinez, D. (2014b). On the measurement of strength of preference in units of money. *Economic Record*, 90, 1–15.

- Calvó-Armengol, A. and Jackson, M. O. (2004). The effects of social networks on employment and inequality. *American Economic Review*, 94(3), 426–454.
- Cerreia-Vioglio, S., Dillenberger, D. and Ortoleva, P. (2015). Cautious expected utility and the certainty effect. *Econometrica*, 83(2), 693–728.
- Coombs, C. and Huang, L. (1976). Tests of the betweenness property of expected utility. *Journal of Mathematical Psychology*, 13, 323–337.
- Cox, J. C. and Sadiraj, V. (2006). Small- and large-stakes risk aversion: Implications of concavity calibration for decision theory. *Games and Economic Behavior*, 56(1), 45–60.
- Dominitz, J. and Manski, C. F. (2004). How should we measure consumer confidence? *Journal of Economic Perspectives*, 18(2), 51–66.
- Elliott, M., Golub, B. and Jackson, M. O. (2014). Financial networks and contagion. *American Economic Review*, 104(10), 3115–3153.
- Elster, J. (1987). *The Multiple Self*. Cambridge University Press: Cambridge.
- Fechner, G. T. (1860/1966). *Elements of psychophysics*. New York: Holt, Rinehart and Winston.
- Gigerenzer, G. and Selten, R. (2001). *Bounded Rationality: The Adaptive Toolbox*. MIT Press.
- Gilboa, I. and Schmeidler, D. (1995). Case-based decision theory. *Quarterly Journal of Economics*, 110 (3), 605–639.
- Gold, J. I. and Shadlen, M. N. (2007). The neural basis of decision making. *Annual Review of Neuroscience*, 30, 535–574.
- Gouriéroux, C. and Monfort, A. (1997). *Simulation-based Econometric Methods*. Oxford University Press: Oxford.
- Hammond, P. (1993). Interpersonal comparisons of utility: Why and how they are and should be made. In J. Elster and J. Roemer (Eds.), *Interpersonal Comparisons of Well-Being*. Cambridge University Press: London.
- Harless, D. W. and Camerer, C. F. (1994). The predictive utility of generalized expected utility theories. *Econometrica*, 62, 1251–1289.
- Harstad, R. M. and Selten, R. (2013). Bounded-rationality models: Tasks to become intellectually competitive. *Journal of Economic Literature*, 51(2), 496–511.
- Hausmann, D. and Läge, D. (2008). Sequential evidence accumulation in decision making: The individual desired level of confidence can explain the extent of information acquisition. *Judgment and Decision Making*, 3(3), 229–43.
- Hey, J. D. and Orme, C. (1994). Investigating generalizations of expected utility theory using experimental data. *Econometrica*, 62, 1291–1326.
- Holt, C. A. and Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644–1655.
- Jamieson, D. G. and Petrusic, W. M. (1977). Preference and time to choose. *Organizational Behavior and Human Performance*, 19, 56–67.
- Johnson, J. G. and Busemeyer, J. R. (2005). A dynamic, computational model of preference reversal phenomena. *Psychological Review*, 112(4), 841–861.
- Johnson, E. J., Häubl, G. and Keinan, A. (2007). Aspects of endowment: A query theory of value construction. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 33, 461–474.
- Kahneman, D. and Tversky, A. (1979). Prospect Theory: An analysis of decision under risk. *Econometrica*, 47, 263–291.
- Kahneman, D., Rosenfield, A. M, Gandhi, L. and Blaser, T. (2016). Inconsistent decision making is a huge hidden cost for many companies. Here’s how to overcome what we call noise. *Harvard Business Review*, 94(10), 38–46.

- Kocher, M., Pahlke, J. and Trautmann, S. T. (2013). *Tempus fugit*: Time pressure in risky decisions. *Management Science*, 59(10), 2380–2391.
- Loomes, G. (2014). Quantitative tests of the Perceived Relative Argument Model: reply to Guo and Regenwetter. *Psychological Review*, 121, 706–10.
- Luce, R. D. and Suppes, P. (1965). Preferences, utility, and subjective probability. In R. D. Luce *et al.* (Eds.), *Handbook of mathematical psychology*, 3, New York: John Wiley.
- Loomes, G., Moffatt, P. G. and Sugden, R. (2002). A Microeconomic test of alternative theories of risky choice. *Journal of Risk and Uncertainty*, 24, 103–130.
- Luce, R. D. (1992). A theory of certainty equivalents for uncertain alternatives. *Journal of Behavioral Decision Making*, 5, 201–216.
- Luce, R. D., Mellers, B. and Chang, S.J. (1993). Is choice the correct primitive? On using certainty equivalents and reference levels to predict choices among gambles. *Journal of Risk and Uncertainty*, 6, 115–143.
- Ludvigson, S. C. (2004). Consumer confidence and consumer spending. *Journal of Economic Perspectives*, 18(2), 29–50.
- Machina, M. (1982). ‘Expected utility’ analysis without the independence axiom. *Econometrica*, 50, 277–323.
- Marschak, J. (1950). Rational behavior, uncertain prospects, and measurable utility. *Econometrica*, 18(2), 111–141.
- Moffatt, P. (2005). Stochastic choice and the allocation of cognitive effort. *Experimental Economics*, 8, 369–388.
- Moffatt, P. and Peters, S. (2001). Testing for the presence of a tremble in economic experiments. *Experimental Economics*, 4, 221–228.
- Mosteller, F. and Nogee, P. (1951). An experimental measurement of utility. *Journal of Political Economy*, 59, 371–404.
- Munier, B., Selten, R., Bouyssou, D., Bourguine, P. and Day, R. (1999). Bounded rationality modeling. *Marketing Letters*, 10(3), 233–248.
- Myers, J., Suydam, M. and Gambino, B. (1965). Contingent gains and losses in a risk-taking situation. *Journal of Mathematical Psychology*, 2, 363–70.
- Otter, T., Johnson, J., Rieskamp, J., Allenby, G. M., Brazell, J. D., Diedrich, A., Hutchinson, J. W., MacEachern, S., Ruan, S. and Townsend, J. (2008). Sequential sampling models of choice: Some recent advances. *Marketing Letters*, 19, 255–267.
- Palacios-Huerta, I. and Serrano, R. (2006). Rejecting gambles under expected utility. *Economics Letters*, 91, 250–259.
- Payne, J. W., Bettman, J. R. and Johnson, E. J. (1993). *The Adaptive Decision Maker*. Cambridge: Cambridge University Press.
- Petrusic, W. M. and Jamieson, D. G. (1978). Relation between probability of preferential choice and time to choose changes with practice. *Journal of Experimental Psychology: Human Perception and Performance*, 4, 471–482.
- Pleskac, T. J. and Busemeyer, J. R. (2010). Two-stage dynamic signal detection: A theory of choice, decision time, and confidence. *Psychological Review*, 117, 864–901.
- Rabin, M. (2000). Risk aversion and expected-utility theory: A calibration theorem. *Econometrica*, 68(5), 1281–1292.
- Ratcliff, R. and Smith, P.L. (2004). A Comparison of Sequential Sampling Models for Two-Choice Reaction Time. *Psychological Review*, 111, 333–367.
- Reiss, J. (2011). A plea for (good) simulations: Nudging economics towards an experimental science. *Simulation & Gaming*, 42(2), 243–264.
- Rieskamp, J., Busemeyer, J. and Mellers, B. (2006). Extending the bounds of rationality: Evidence and theories of preferential choice. *Journal of Economic Literature*, 44, 631–661.

- Roe, R. M., Busemeyer, J. R. and Townsend, J. T. (2001). Multi-alternative decision field theory: A dynamic connectionist model of decision making. *Psychological Review*, 108, 370–392.
- Rubinstein, A. (1988). Similarity and decision making under risk: is there a utility theory resolution to the Allais paradox? *Journal of Economic Theory*, 46, 145–153.
- Rubinstein, A. (2001). Comments on the risk and time preferences in economics. Unpublished manuscript, available at http://www.econ.nyu.edu/user/bisina/rubinstein_onRabin.pdf.
- Rubinstein, A. (2007). Instinctive and cognitive reasoning: a study of response times. *The Economic Journal*, 117, 1243–1259.
- Rubinstein, A. (2013). Response time and decision making: an experimental study. *Judgment and Decision Making*, 8(5), 540–551.
- Rubinstein, A. (2016). *Lecture Notes in Microeconomic Theory*. Princeton University Press: Princeton and Oxford.
- Shiryayev, A. N. (1978). *Optimal Stopping Rules*. Springer-Verlag: Berlin Heidelberg.
- Sieck, W. and Yates, J.F. (1997). Exposition effects on decision making: Choice and confidence in choice. *Organizational Behavior and Human Decision Processes*, 70, 207–219.
- Simon, H. A. (1978). Rationality as process and as product of thought. *American Economic Review, Papers and Proceedings*, 68(2), 1–16.
- Spiliopoulos, L. and Ortmann, A. (2016). The BCD of response time analysis in experimental economics. Available at SSRN: <http://ssrn.com/abstract=2401325>.
- 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., Chater, N. and Brown, G. D. A. (2006). Decision by sampling. *Cognitive Psychology*, 53, 1–26.
- Stewart, N., Reimers, S. and Harris, A. J. L. (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.
- Stigler, G. J. (1961). The economics of information. *Journal of Political Economy*, 69(3), 213–225.
- Stott, H.P. (2006). Cumulative prospect theory's functional menagerie. *Journal of Risk and Uncertainty*, 32, 101–130.
- Tversky, A. (1972). Elimination by aspects: A theory of choice. *Psychological Review*, 79(4), 281–299.
- Tversky, A. and Russo, J. E. (1969). Substitutability and similarity in binary choices. *Journal of Mathematical Psychology*, 6, 1–12.
- Tyebjee, T. (1979). Response time, conflict and involvement in brand choice. *Journal of Consumer Research*, 6, 295–304.
- Wilcox, N. T. (1993). Lottery choice: Incentives, complexity and decision time. *The Economic Journal*, 103(421), 1397–1417.
- Wilcox, N. T. (2008). Stochastic models for discrete choice under risk: A critical primer and econometric comparison. In J. C. Cox and G. W. Harrison (Eds.), *Research in experimental economics*, 12. Bingley: Emerald Group Publishing Limited.
- Wilcox, N. T. (2011). 'Stochastically more risk averse': A contextual theory of stochastic discrete choice under risk. *Journal of Econometrics*, 162, 89–104.
- Wilcox, N. T. (2016). Random expected utility and certainty equivalents: Mimicry of probability weighting functions. MPRA Paper No. 73173.




FIGURE 1
 Lotteries in the Marschak-Machina triangle




FIGURE 2
 A sketch of BREUT's indifference curves for the limiting case ($d \rightarrow 0$)

TABLE 1

Choice between a fixed lottery $B = (\mathbf{40}, 0.8; \mathbf{0}, 0.2)$ and increasing sure amounts of money A
 $(\alpha = 0.35, \beta = 1.0, d = 0.1)$

A	Pr Core ($A \succ B$)	$E[V(A, B)]$	$\Pr(A \succ B)$	RT	$Conf^*$
(20 , 1)	0.04	-7.58	0.00	7.02	0.87
(22 , 1)	0.08	-5.60	0.00	7.56	0.85
(24 , 1)	0.15	-3.60	0.02	8.46	0.82
(26 , 1)	0.26	-1.58	0.13	9.90	0.77
(28 , 1)	0.45	0.42	0.50	10.99	0.73
(30 , 1)	0.73	2.41	0.92	9.63	0.78
(32 , 1)	0.97	4.41	1.00	7.83	0.84

TABLE 2

Changing the median risk aversion parameter α ($\beta = 1.0, d = 0.1$)

$A = (\mathbf{30}, 1)$
 $B = (\mathbf{40}, 0.8; \mathbf{0}, 0.2)$

α	Pr Core ($A \succ B$)	$E[V(A, B)]$	$\Pr(A \succ B)$	RT	$Conf^*$
0.05	0.20	-1.36	0.05	9.12	0.80
0.10	0.28	-0.87	0.13	10.01	0.77
0.15	0.36	-0.39	0.28	10.73	0.74
0.20	0.45	0.19	0.49	11.01	0.73
0.25	0.55	0.84	0.68	10.84	0.74
0.30	0.64	1.57	0.83	10.32	0.76
0.35	0.73	2.41	0.92	9.65	0.78
0.40	0.80	3.39	0.97	9.01	0.80

TABLE 3

Changing the range of the distribution of risk aversion coefficients β ($\alpha = 0.30, d = 0.1$)

$$A = (30, 1)$$

$$B = (40, 0.8; 0, 0.2)$$

β	Pr Core ($A \succ B$)	$E[V(A, B)]$	$\Pr(A \succ B)$	RT	$Conf^*$
0.25	0.94	0.96	1.00	7.72	0.84
0.40	0.82	1.01	0.97	8.81	0.81
0.55	0.75	1.10	0.93	9.51	0.78
0.70	0.70	1.21	0.89	9.97	0.77
0.85	0.66	1.38	0.86	10.18	0.76
1.00	0.64	1.60	0.83	10.31	0.76
1.15	0.62	1.78	0.81	10.38	0.75
1.30	0.61	2.13	0.80	10.43	0.75

TABLE 4

Changing the desired level of confidence decrease parameter d ($\alpha = 0.30, \beta = 1.0$)

$$A = (30, 1)$$

$$B = (40, 0.8; 0, 0.2)$$

d	Pr Core ($A \succ B$)	$E[V(A, B)]$	$\Pr(A \succ B)$	RT	$Conf^*$
0.50	0.64	1.57	0.75	6.00	0.50
0.40	0.64	1.57	0.76	6.42	0.59
0.30	0.64	1.57	0.77	7.01	0.63
0.20	0.64	1.57	0.79	8.00	0.68
0.15	0.64	1.57	0.81	8.81	0.71
0.10	0.64	1.57	0.84	10.31	0.76
0.05	0.64	1.57	0.89	13.94	0.82
0.03	0.64	1.57	0.92	17.54	0.85
0.02	0.64	1.57	0.94	20.87	0.88
0.01	0.64	1.57	0.96	27.75	0.92

TABLE 5

Choosing between dominating and dominated lotteries ($\alpha = 0.23, \beta = 1.0, d = 0.1$)

<i>A</i> <i>Dominating</i>	<i>B</i> <i>Dominated</i>	Pr Core ($A \succ B$)	$E[V(A, B)]$	$\Pr(A \succ B)$	<i>RT</i>	<i>Conf*</i>
(60, 0.5; 0, 0.5)	(50, 0.5; 0, 0.5)	1.00	3.50	1.00	8.36	0.89
(51, 0.5; 0, 0.5)	(50, 0.5; 0, 0.5)	1.00	0.35	1.00	8.35	0.89

TABLE 6

Comparing $K = (180, 0.25; 0, 0.75)$ and $L = (40, 0.8; 0, 0.2)$ against sure amounts M from 25 to 33($\alpha = 0.23, \beta = 1.0, d = 0.1$)

<i>M</i>	$K = (180, 0.25; 0, 0.75)$				
	Pr Core ($K \succ M$)	$E[V(K, M)]$	$\Pr(K \succ M)$	<i>RT</i>	<i>Conf*</i>
(25, 1)	0.63	4.65	0.76	10.68	0.74
(26, 1)	0.60	3.55	0.71	10.80	0.74
(27, 1)	0.57	2.51	0.66	11.01	0.73
(28, 1)	0.55	1.56	0.61	11.07	0.73
(29, 1)	0.52	0.57	0.54	11.12	0.73
(30, 1)	0.49	-0.43	0.48	11.10	0.73
(31, 1)	0.47	-1.43	0.42	11.12	0.73
(32, 1)	0.44	-2.40	0.35	10.97	0.73
(33, 1)	0.41	-3.36	0.30	10.87	0.74

<i>M</i>	$L = (40, 0.8; 0, 0.2)$				
	Pr Core ($L \succ M$)	$E[V(L, M)]$	$\Pr(L \succ M)$	<i>RT</i>	<i>Conf*</i>
(25, 1)	0.94	4.43	1.00	7.36	0.85
(26, 1)	0.90	3.45	0.99	7.86	0.84
(27, 1)	0.84	2.44	0.97	8.66	0.81
(28, 1)	0.76	1.44	0.91	9.68	0.78
(29, 1)	0.64	0.44	0.71	10.71	0.74
(30, 1)	0.49	-0.56	0.39	10.95	0.74
(31, 1)	0.31	-1.56	0.11	9.93	0.77
(32, 1)	0.13	-2.56	0.01	8.45	0.82
(33, 1)	0.01	-3.56	0.00	7.66	0.84

TABLE 7

BREUT's predictions for the independence and betweenness pairs of Figure 2a ($\alpha = 0.23, \beta = 1.0$)

$$\begin{aligned}
A &= (\mathbf{30}, 1) \\
B &= (\mathbf{40}, 0.8; \mathbf{0}, 0.2) \\
C &= (\mathbf{30}, 0.25; \mathbf{0}, 0.75) \\
D &= (\mathbf{40}, 0.2; \mathbf{0}, 0.8) \\
E &= (\mathbf{40}, 0.2; \mathbf{30}, 0.75; \mathbf{0}, 0.05)
\end{aligned}$$

d	Safer (S)	Riskier (R)	Pr Core ($S \succ R$)	$E[V(S, R)]$	$\Pr(S \succ R)$	RT	$Conf^*$
0.10	A	B	0.51	0.55	0.61	10.97	0.72
	C	D	0.51	-0.07	0.41	14.42	0.76
	A	E	0.51	0.15	0.62	14.64	0.72
	E	B	0.51	0.41	0.61	18.33	0.72
0.05	A	B	0.51	0.55	0.65	15.94	0.78
	C	D	0.51	-0.07	0.37	20.86	0.82
	A	E	0.51	0.15	0.65	21.18	0.78
	E	B	0.51	0.41	0.65	26.71	0.77
0.01	A	B	0.51	0.55	0.80	44.34	0.85
	C	D	0.51	-0.07	0.23	59.91	0.91
	A	E	0.51	0.15	0.81	58.31	0.86
	E	B	0.51	0.41	0.80	73.92	0.85

TABLE 8

BREUT's predictions for alternative betweenness pairs ($d = 0.01$)

$$\begin{aligned}
A' &= (\mathbf{25}, 1) \\
A'' &= (\mathbf{35}, 1) \\
B &= (\mathbf{40}, 0.8; \mathbf{0}, 0.2) \\
E &= (\mathbf{40}, 0.2; \mathbf{30}, 0.75; \mathbf{0}, 0.05)
\end{aligned}$$

Core r values	Safer (S)	Riskier (R)	Pr Core ($S \succ R$)	$E[V(S, R)]$	$\Pr(S \succ R)$	RT	$Conf^*$
0 and 0.7205492791	A'	B	0.50	0.00	0.50	34.39	0.90
	E	B	0.50	-0.09	0.48	56.14	0.90
0 and -3.3503455478	A''	B	0.50	0.00	0.50	34.64	0.90
	E	B	0.50	0.05	0.53	54.93	0.90