Dynamic Choice Behavior in a Natural Experiment

by

Steffen Andersen, Glenn W. Harrison, Morten Igel Lau and Elisabet E. Rutström †

October 2006

Abstract. We examine dynamic choice behavior in a natural experiment with large stakes and a demographically divers sample. The television game show *Deal Or No Deal* offers a rich paradigm to examine the latent decision processes that people use to make choices under uncertainty when they face future options linked to current choices. We have three major findings. First, we show that popular utility functions that assume constant relative or absolute risk aversion and expected utility theory defined over the prizes cannot characterize these choices, which exhibit increasing relative risk aversion over prizes ranging from a penny to nearly half a million U.S. dollars. Second, the argument of the utility function under expected utility theory reflects the integration of game show prizes with regular income. These decision makers do not segregate the income from the lotteries they face on the game show from the income that they bring to the game show. Allowing for this integration of income and game show prizes leads to choice behavior consistent with constant relative risk aversion. Third, we examine the effects of allowing contestants to make choices characterized by non-standard decision models. We find evidence of some probability weighting, but no loss aversion. Hence we identify two senses in which this large-stakes, naturally occurring behavior differs from behavior characterized in a small-stakes, laboratory setting: individuals appear to integrate prizes with some broader measure of baseline income or consumption, and individuals do not appear to exhibit systematic loss aversion.

† Centre for Economic and Business Research, Copenhagen Business School, Copenhagen, Denmark (Andersen); Department of Economics, College of Business Administration, University of Central Florida, USA (Harrison and Rutström) and Department of Economics and Finance, Durham Business School, Durham University, United Kingdom (Lau). E-mail: SA.CEBR@CBS.DK, GHARRISON@RESEARCH.BUS.UCF.EDU, M.I.LAU@DURHAM.AC.UK, and ERUTSTROM@BUS.UCF.EDU.
Harrison and Rutström thank the U.S. National Science Foundation for research support under grants NSF/IIS 9817518, NSF/HSD 0527675 and NSF/SES 0616746. We are grateful to Hans Jørgen Andersen, Hazel Harrison, William Harrison and Chris Powney for assistance collecting data, and to Pavlo Blavatskyy, Ganna Pogrebna, Thierry Post and Martijn van den Assem for discussions and comments.

The television game show *Deal Or No Deal* provides a wonderful opportunity to examine dynamic choice under uncertainty in a controlled manner with substantial stakes. Observed behavior in these shows constitutes a controlled natural experiment: contestants are presented with well-defined dynamic choices where the stakes are real and sizeable, and the tasks are repeated in the same manner from contestant to contestant.[1]

The game involves each contestant deciding in a given round whether to accept a deterministic cash offer or to continue to play the game. It therefore represents a non-strategic game of timing, and is often presented to contestants as exactly that by the host. If the subject chooses "No Deal," and continues to play the game, then the outcome is uncertain. The sequence of choices is intrinsically dynamic because the deterministic cash offer evolves in a relatively simple manner as time goes on. Apart from adding drama to the show, this temporal connection makes the choices particularly interesting and, arguably, more relevant to the types of decisions one expects in naturally occurring environments.[2] We explain the format of the show in section 1, and discuss this temporal connection.

We examine two general issues in the specification of dynamic choice behavior.

The first issue is the *characterization of this behavior assuming expected utility theory* (EUT). This domain presents a striking challenge to many of the specifications of EUT that are often applied in laboratory settings: the stakes are huge and span a wide range. Echoing findings by Holt and Laury [2002], we show that one must abandon restrictive functional forms that assume constant relative or

[1] Game shows are increasingly recognized as a valuable source of replicable data with large stakes. For example, see Beetsma and Schotman [2001], Berk, Hughson and Vandezande [1996], Février and Linnemer, [2006], Gertner [1993], Hartley, Lanot and Walker [2005], Healy and Noussair [2004], Levitt [2004], List [2006], Metrick [1995], and Tenorio and Cason [2002]. Because the task domain of *Deal Or No Deal* is so rich, it has attracted a large number of complementary analyses. In an appendix (available on request) we compare our formal analysis to those of Blavatskyy and Pogrebna [2006c], Bombardini and Trebbi [2005], Botti, Conte, Di Cagno and D'Ippoliti. [2006], Deck, Jungmin and Reyes [2006], Mulino, Schelings, Brooks and Faff [2006], Post, van den Assem, Baltussen and Thaler [2006] and De Roos and Sarafidis [2006].

[2] Cubitt and Sugden [2001] make this point explicitly, contrasting the static, one-shot nature of the choice tasks typically encountered in laboratory experiments with the sequential, dynamic choices that theory is supposed to be applied to in the field. It is also clearly stated in Thaler and Johnson [1990; p. 643], who recognize that the issues raised by considering dynamic sequences of choices are "quite general since decisions are rarely made in temporal isolation."

absolute risk aversion when utility is defined over such large prizes. Consistent with the perspectives of Rubinstein [2002] and Cox and Sadiraj [2006], responding to the criticisms of EUT by Rabin [2000], we also find that one must also allow some flexibility about the proper arguments of the utility function. The assumption that utility is defined solely over prizes must be abandoned in favor of the assumption that the utility of prizes in the game show is integrated with some measure of income. On the other hand, allowing for utility functions in which prizes and income are integrated leads to choices consistent with constant relative risk aversion.

The second issue is the characterization of behavior using alternatives to EUT, particularly alternatives that allow for sign and rank-dependent preferences. The *Deal Or No Deal* (*DOND*) environment would seem to be a natural breeding ground for these behaviors, since a key feature of the game is the strikingly asymmetric distribution of prizes. We find evidence that there is indeed some probability weighting being undertaken by contestants, particularly in the gain domain. But we find *no evidence of loss aversion* using a natural assumption about the reference point determining gains and losses.

The structure of this game provides a canonical test-bed for many of the issues that arise when characterizing dynamic choice behavior. We are not just interested in figuring out if contestants play this specific game consistently with one theory or another, or what risk attitudes or probability weighting subjects employ. Those are worthwhile exercises, given the stakes involved here and the relatively divers demographic mix in the sample. But this framework forces one to be explicit about issues that can probably be safely ignored in static choice experiments in a standard laboratory environment, particularly the range of the stakes and the effects of path dependence.

One feature of our approach to these data is that we write out explicit models of the latent structure underlying the observed choices, so that one can directly estimate parameters of the theory providing that latent structure. This does require us to impose some parametric structure on these theories, but we use flexible functional forms that have been popular in other applications. One benefit if this approach is that we explicitly account for uncertainty in estimates of key parameters, so that overall inferences about model performance can be conditioned on that uncertainty.

Another feature of our approach is that we treat EUT and alternative models symmetrically. We examine the specification issues that arise when one assumes that the observed data are characterized solely by one or the other model. However, we reject the view that one theory is always right (or wrong), or even that one theory is necessarily the correct theory for a given domain.

In section 1 we document the game show format and field data we use. We employ data from the version of the show in the United Kingdom, reflecting 1,074 choices by 211 contestants over prizes ranging from 1 penny to £250,000. This prize range is roughly equivalent to US $0.02 and US $460,000. Average earnings in the game show are £16,750 in our sample. The distribution of earnings is heavily skewed, with relatively few subjects receiving the highest prizes, and median earnings are £13,000. In section 2 we describe the general statistical models developed for these data, assuming an EUT model of the latent decision-making process, and present results conditional on those assumptions. In section 3 we consider a number of extensions of the core analysis, going beyond the assumption that one EUT specification explains all choices. Section 4 briefly reviews major, open issues that are motivated by our findings and this extraordinarily rich task environment. Finally, section 5 draws some conclusions.

1. The Naturally Occurring Game Show Data

The version of *Deal Or No Deal* shown in the United Kingdom starts with a contestant being randomly picked from a group of 22 preselected people. They are told that a known list of monetary prizes, ranging from 1p up to £250,000, has been placed in 22 boxes. Each box has a number from 1 to 22 associated with it, and one box has been allocated at random to the contestant before the show. The contestant is informed that the money has been put in the box by an independent third party, and in fact it is common that any unopened boxes at the end of play are opened so that the audience can see that all prizes were in play. The picture below shows how the prizes are displayed to the subject, the proto-typically British "Trevor," at the beginning of the game.

In round 1 the contestant must pick 5 of the remaining 21 boxes to be opened, so that their prizes can be displayed. A good round for a contestant occurs if the opened prizes are low, and hence the odds increase that his box holds the higher prizes. At the end of each round the host is phoned by a "banker" who makes a deterministic cash offer to the contestant.

The initial offer in early round is typically low in comparison to expected offers in later rounds. We document an empirical offer function later, but the qualitative trend is quite clear: the bank offer starts out at roughly 15% of the expected value of the unopened boxes, and increases to roughly 24%, 34%, 42%, 54% and then 73% in rounds 2 though 6. This trend is significant, and serves to keep all but extremely risk averse contestants in the game for several rounds. For this reason it is clear that the box that the contestant "owns" has an option value in future rounds.

In round 2 the contestant must pick 3 boxes to open, and then there is another bank offer to consider. In rounds 3 through 6 the contestant must open 3 boxes in each round. At the end of round 6 there are only 2 unopened boxes, one of which is the contestant's box.[3]

In round 6 the decision is a relatively simple one from an analyst's perspective: either take the non-stochastic cash offer or take the lottery with a 50% chance of either of the two remaining unopened prizes. We could assume some latent utility function, or non-standard decision function, and directly estimate parameters for that function that best explains the observed binary choices in this round. Unfortunately, relatively few contestants get to this stage, having accepted offers in

[3] Some versions substitute the option of switching the contestant's box for an unopened box, instead of a bank offer. This is particularly common in the French and Italian versions. The UK version does not generally have this feature, making the analysis much cleaner. In our UK sample it only occurred for 3 subjects in round 1.

earlier rounds. In our data, only 39% of contestants reach that point.[4] More serious than the smaller sample size, one naturally expects that risk attitudes would affect those surviving to this round. Thus there would be a serious sample selection bias if one just studied choices in later rounds.

In round 5 the decision is conceptually much more interesting. Again the contestant can just take the non-stochastic cash offer. But now the decision to continue amounts to opting for one of two potential lotteries: (i) take the offer that will come in round 6 after one more box is opened, or (ii) decide in round 5 to reject that offer, and then play out the final 50/50 lottery. Each of these is an uncertain lottery, from the perspective of the contestant in round 5. Choices in earlier rounds involve larger and larger sets of potential lotteries of this form.

If the bank offer was random this sequence could be evaluated as a series of static choices, at least under standard EUT. The cognitive complexity of evaluating the compound lottery might be a factor in behavior, but it would conceptually be simple to analyze. However, the bank offer gets richer and richer over time, *ceteris paribus* the random realizations of opened boxes. In other words, if each unopened box truly has the same subjective probability of having any remaining prize, there is a positive expected return to staying in the game for more and more rounds. Thus a risk averse subject that might be just willing to accept the bank offer, *if the offer were not expected to get better and better*, would choose to continue to another round since the expected improvement in the bank offer provides some compensation for the additional risk of going into the another round.

Thus, to evaluate the parameters of some latent utility function given observed choices in earlier rounds, we have to mentally play out all possible future paths that the contestant faces.[5] Specifically, we have to play out those paths assuming the values for the parameters of the likelihood function, since they affect when the contestant will decide to "Deal" with the banker, and hence the

---

[4] This fraction is even smaller in other versions of the game show in other countries, where there are typically 9 rounds. Other versions generally have bank offers that are more generous in later rounds, with most of them approaching 100% of the expected value of the unopened boxes. In some cases the offers exceed 100% of this expected value.

[5] Or make some *a priori* judgement about the bounded rationality of contestants. For example, one could assume that contestants only look forward one or two rounds, or that they completely ignore bank offers.

expected utility of the compound lottery. This corresponds to procedures developed in the finance literature to price path-dependant derivative securities using Monte Carlo simulation (e.g., Campbell, Lo and MacKinlay [1997; §9.4]).

Saying "No Deal" in early rounds provides one with the option of being offered a better deal in the future, *ceteris paribus* the expected value of the unopened prizes in future rounds. Since the process of opening boxes is a martingale process, even if the contestant gets to pick the boxes to be opened, it has a constant *future* expected value in any given round equal to the *current* expected value. This implies, given the exogenous bank offers (as a function of expected value),[6] that the dollar value of the offer will get richer and richer as time progresses. Thus bank offers themselves will be a sub-martingale process.

The show began broadcasting in the United Kingdom in October 2005, and has been showing constantly since. There are normally 6 episodes per week: a daytime episode and a single prime time episode, each roughly 45 minutes in length. Our data are drawn primarily from direct observation of recorded episodes, but we also verify data against those tabulated on the web site http://www.dond.co.uk/.

2. Estimation Assuming EUT

We evaluate behavior sequentially, starting with a simple set of assumptions about the way in which the observed choices are made and then adding variations. To start we assume that subjects follow EUT and that they use a utility function defined over the prizes on offer. In this section we explain the basic logic of our estimation approach, and present it formally.[7] In section 3 we consider

[6] Things become much more complex if the bank offer in any round is statistically informative about the prize in the contestant's box. In that case the contestant has to make some correction for this possibility, and also consider the strategic behavior of the banker's offer. Bombardini and Trebbi [2005] offer clear evidence that this occurs in the Italian version of the show, but there is no evidence that it occurs in the U.K. version.

[7] Bombardini and Trebbi [2005], Mulino, Schelings, Brooks and Faff [2006] and De Roos and Sarafidis [2006] present formal statistical models that have many of the same features. An appendix (available on request) discusses the differences and similarities. In brief, (a) our model considers more flexible functional forms under EUT; (b) estimates a Prospect Theory specification with allowance for probability weighting and loss aversion, rather than a rank-dependant utility model that only allows for probability weighting; and (c) estimates a psychological model that explicitly allows for contestants to balance aspiration thresholds with

a number of variants of this approach.

A. Basic Intuition

The basic logic of our approach can be explained from the data and simulations shown in Table 1. There are 6 rounds in which the banker makes an offer, and in round 7 the surviving contestant simply opens his box. We observed 211 contestants play the game. Only 45, or 21%, made it to round 7, with most accepting the banker's offer in rounds 4, 5 and 6. The average offer is shown in column 4. We stress that this offer is stochastic from the perspective of the sample as a whole, even if it is non-stochastic to the specific contestant in that round. Thus, to see the logic of our approach from the perspective of the individual decision-maker, think of the offer as a non-stochastic number, using the average values shown as a proximate indicator of the value of that number in a particular instance.

In round 1 the contestant might consider up to 6 *virtual lotteries*. He might look ahead one round and contemplate the outcomes he would get if he turned down the offer in round 1 and accepted the offer in round 2. This virtual lottery, realized in virtual round 2 in the contestant's thought experiment, would generate an average payoff of £7,422 with a standard deviation of £7,026. The top panel of Figure 1 shows the simulated distribution of this particular lottery. The distribution of payoffs to these virtual lotteries are highly skewed, so the standard deviation may be slightly misleading if one thinks of these as Gaussian distributions. However, we just use the standard deviation as one pedagogic indicator of the uncertainty of the payoff in the virtual lottery: in our formal analysis we consider the complete distribution of the virtual lottery in a non-

---

utility evaluation. Each of these extensions make significant differences to the analysis. Unlike Bombardini and Trebbi [2005] we correct for the fact that each contestant offers several responses that might be correlated with unobservables. Blavatskyy and Pogerbna [2006c] and Botti, Conte, DiCagno and D'Ippoliti [2006] estimate models that do not allow for the "continuation value" of saying "No Deal" in a round: they assume that the contestant compares the current bank offer with the terminal lottery that assumes that they would always say "No Deal" in future rounds. Blavatskyy and Pogerbna [2006c] focus on the role of different stochastic error stories. Botti, Conte, DiCagno and D'Ippoliti [2006] focus on the role of unobserved individual heterogeneity, much like De Roos and Sarafidis [2006]. This approach is particularly valuable for a naturally occurring data set, such as *DOND*, where there are relatively few observable individual characteristics.

parametric manner.

In round 1 the contestant can also consider what would happen if he turned down offers in rounds 1 and 2, and accepted the offer in round 3. This virtual lottery would generate, from the perspective of round 1, an average payoff of £9,704 with a standard deviation of £9,141. The bottom panel of Figure 1 shows the simulated distribution of this particular virtual lottery. Compared to the virtual lottery in which the contestant said "No Deal" in round 1 and "Deal" in round 2, shown above it in Figure 1, it gives less weight to the smallest prizes and greater weight to higher prizes. Similarly for each of the other virtual lotteries shown in Table 1.

The forward looking contestant in round 1 is assumed to behave as if he maximizes the expected utility of accepting the current offer or continuing. The expected utility of continuing, in turn, is given by simply evaluating each of the 6 virtual lotteries shown in the first row of Table 1. The average payoff increases steadily, but so does the standard deviation of payoffs, so this evaluation requires knowledge of the utility function of the contestant. Given that utility function, the contestant is assumed to behave as if they evaluate the expected utility of each of the 6 virtual lotteries. Thus we calculate six expected utility numbers, conditional on the specification of the parameters of the assumed utility function and the virtual lotteries that each subject faces in their round 1 choices. In round 1 the subject then simply compares the *maximum* of these 6 expected utility numbers to the utility of the non-stochastic offer in round 1. If that maximum exceeds the utility of the offer, he turns down the offer; otherwise he accepts it.

In round 2 a similar process occurs. One critical feature of our virtual lottery simulations is that they are conditioned on the actual outcomes that each contestant has faced in prior rounds. Thus, if a (real) contestant has tragically opened up the 5 top prizes in round 1, that contestant would not see virtual lotteries such as the ones in Table 1 for round 2. They would be conditioned on that player's history in round 1. We report here averages over all players and all simulations. We undertake 10,000 simulations for *each* player in *each* round, so as to condition on their history.[8]

[8] If bank offers were a deterministic and known function of the expected value of unopened prizes, we would not need anything like 10,000 simulations for later rounds. For the last few rounds of a full game, in

This example can also be used to illustrate how our maximum likelihood estimation procedure works. Assume some specific utility function and some parameter values for that utility function. The utility of the non-stochastic bank offer in round R is then directly evaluated. Similarly, the virtual lotteries in each round R can then be evaluated.[9] They are represented numerically as 20-point discrete approximations, with 20 prizes and 20 probabilities associated with those prizes. Thus, by implicitly picking a virtual lottery over an offer, it is as if the subject is taking a draw from this 20-point distribution of prizes. In fact, they are playing out the *DOND* game, but this representation as a virtual lottery draw is formally identical. The evaluation of these virtual lotteries generates v(R) expected utilities, where v(1)=6, v(2)=5,...,v(6)=1 as shown in Table 1. The maximum expected utility of these v(R) in a given round R is then compared to the utility of the offer, and the likelihood evaluated in the usual manner.[10]

We present a formal statement of the latent EUT process leading to a likelihood defined over parameters and the observed choices, and then discuss how this intuition changes when we assume alternative, non-EUT processes.

B. Formal Specification

We assume that utility is defined over money m using an Expo-Power (EP) function

$$u(m) = [1-\exp(-\alpha m^{1-r})]/\alpha \tag{1}$$

where α and r are parameters to be estimated, and non-satiation requires $r<1$. This functional form

which the bank offer is *relatively* predictable, the use of this many simulations is a numerically costless redundancy.

[9] There is no need to know risk attitudes, or other preferences, when the distributions of the virtual lotteries are generated by simulation. But there is definitely a need to know these preferences when the virtual lotteries are *evaluated*. Keeping these computational steps separate is essential for computational efficiency, and is the same procedurally as pre-generating "smart" Halton sequences of uniform deviates for later, repeated use within a maximum simulated likelihood evaluator (e.g., Train [2003; p. 224ff.]).

[10] The only complication from using a 20-point approximation might occur when one undertakes probability weighting. However, if one uses rank-dependant probability weighting this issue disappears. For example, a 4-point virtual lottery with prizes 100, 100, 200 and 200, each occurring with probability ¼, is the same as a lottery with prizes 100 and 200 each occurring with probability ½. This point is of some importance for our application when one considers the virtual lottery in which the contestant says "No Deal" to every bank offer. In that virtual lottery there are never more than 17 possible outcomes in the UK version, and in round 6 there are exactly 2 possible outcomes.

was introduced by Saha [1993], and played a central role in the characterization of risk attitudes in the laboratory by Holt and Laury [2002]. Relative risk aversion (RRA) for the EP function is given by $r + \alpha(1-r)m^{1-r}$, so RRA varies with income if $\alpha \neq 0$ and r is therefore the RRA for a zero income level. This function nests Constant Absolute Risk Aversion (CARA) utility functions as r tends to 0, but is not defined for α equal to 0.

As α gets closer to 0 the EP function behaves essentially as a Constant Relative Risk Aversion (CRRA) function, even if it is not defined in the limit. The specific CRRA function $u(m) = m^{1-r}/(1-r)$ has been popular in the literature, since it requires only one parameter to be estimated, and when we refer to CRRA we will refer to this functional form.[11] In this case $r \neq 1$ is the RRA coefficient, and $u(m) = \ln(m)$ for $r = 1$. With this parameterization $r = 0$ denotes risk neutral behavior, $r > 0$ denotes risk aversion, and $r < 0$ denotes risk loving.

The CRRA functional form is understandably popular in theoretical and applied work, but it may not be the right functional form for these data. The striking range of prizes in *DOND* provide a unique test-bed for examining the importance of using more flexible functional forms.[12] Holt and Laury [2002] stressed the need to allow for varying RRA in laboratory experiments when the prizes were scaled from a top prize of \$3.85 to \$77, and then to \$192.50 and finally \$346.50. In their data they found statistically significant evidence of increasing RRA over these prize domains, which are substantial for lab experiments but minuscule from the perspective of the *DOND* game shows. Thus, if ever there was a situation in which one might expect CRRA or CARA to fail, it is in the span of prizes used in these natural experiments. Our analysis will allow us to directly test that

[11] Abdellaoui, Barrios and Wakker [2005] offer a one-parameter version of the EP function which exhibits non-constant RRA for empirically plausible parameter values. It does impose some restrictions on the variations in RRA compared to the two-parameter EP function, but is valuable as a parsimonious way to estimate non-CRRA specifications.

[12] In the Epilogue to a book-length review of the economics of risk and time, Gollier [2001; p.424ff.] writes that "It is quite surprising and disappointing to me that almost 40 years after the establishment of the concept of risk aversion by Pratt and Arrow, our profession has not yet been able to attain a consensus about the measurement of risk aversion. Without such a consensus, there is no hope to quantify optimal portfolios, efficient public risk prevention policies, optimal insurance deductibles, and so on. It is vital that we put more effort on research aimed at refining our knowledge about risk aversion. For unclear reasons, this line of research is not in fashion these days, and it is a shame." He also has similar remarks (pp. 425/6) about the long-standing need for empirical evaluations of restrictive functional forms such as CRRA.

expectation.

Probabilities for each outcome k, p_k, are those that are induced by the task, so expected utility is simply the probability weighted utility of each outcome in each lottery. We return to this issue in more detail below, since it relates to the use of virtual lotteries. There were 20 outcomes in each virtual lottery i, so

$$EU_i = \sum_{k=1, 20} [p_k \times u_k]. \tag{2}$$

Of course, we can view the bank offer as being a degenerate lottery.

A simple stochastic specification was used to specify likelihoods conditional on the model. The EU for each lottery pair was calculated for a candidate estimate of the utility function parameters, and the index

$$\nabla EU = (EU_{BO} - EU_L)/\mu \tag{3}$$

calculated, where EU_L is the lottery in the task, EU_{BO} is the degenerate lottery given by the bank offer, and μ is a Fechner noise parameter following Hey and Orme [1994].[13] The index ∇EU is then used to define the cumulative probability of the observed choice to "Deal" using the cumulative standard normal distribution function:

$$G(\nabla EU) = \Phi(\nabla EU). \tag{4}$$

The likelihood, conditional on the EUT model being true and the use of the EP utility function, depends on the estimate of r, α and μ given the above specification and the observed choices. The conditional log-likelihood is

$$\ln L^{EUT}(r, \alpha, \mu; y) = \sum_i [(\ln G(\nabla EU) \mid y_i=1) + (\ln (1-G(\nabla EU)) \mid y_i=0)] \tag{5}$$

where y_i =1(0) denotes the choice of "Deal" ("No Deal") in task i.

We extend this standard formulation to include forward looking behavior by redefining the

[13] Harless and Camerer [1994], Hey and Orme [1994] and Loomes and Sugden [1995] provided the first wave of empirical studies including some formal stochastic specification in the version of EUT tested. There are several species of "errors" in use, reviewed by Hey [1995][2002], Loomes and Sugden [1995], Ballinger and Wilcox [1997], and Loomes, Moffatt and Sugden [2002]. Some place the error at the final choice between one lottery or the other after the subject has decided deterministically which one has the higher expected utility; some place the error earlier, on the comparison of preferences leading to the choice; and some place the error even earlier, on the determination of the expected utility of each lottery.

lottery that the contestant faces. One such virtual lottery reflects the possible outcomes if the subject always says "No Deal" until the end of the game and receives his prize. We call this a virtual lottery since it *need* not happen; it *does* happen in some fraction of cases, and it *could* happen for any subject.

Similarly, we can substitute other virtual lotteries reflecting other possible choices by the contestant. Just before deciding whether to accept the bank offer in round 1, what if the contestant behaves as if the following simulation were repeated Γ times:

> { Play out the remaining 5 rounds and pick boxes at random until all but 2 boxes are unopened. Since this is the last round in which one would receive a bank offer, calculate the expected value of the remaining 2 boxes. Then multiply that expected value by the fraction that the bank is expected to use in round 6 to calculate the offer. Pick that fraction from a prior as to the average offer fraction, recognizing that the offer fraction is stochastic. }

The end result of this simulation is a sequence of Γ virtual bank offers in round 6, viewed from the perspective of round 1. This sequence then defines the virtual lottery to be used for a contestant in round 1 whose horizon is the last round in which the bank will make an offer. Each of the Γ bank offers in this virtual simulation occurs with probability $1/\Gamma$, by construction. To keep things numerically manageable, we can then take a 20-point discrete approximation of this lottery, which will typically consist of Γ distinct real values, where one would like Γ to be relatively large (we use Γ=10,000). This simulation is conditional on the 5 boxes that the subject has already selected at the end of round 1. Thus the lottery reflects the historical fact of the 5 specific boxes that this contestant has already opened.

The same process can be repeated for a virtual lottery that only involves looking forward to the expected offer in round 5. And for a virtual lottery that only involves looking forward to rounds 4, 3 and 2, respectively. Table 1 illustrates the outcome of such calculations. The contestant can be viewed as having a set of 6 virtual lotteries to compare, each of which entail saying "No Deal" in round 1. The different virtual lotteries imply different choices in future rounds, but the same response in round 1.

To decide whether to accept the deal in round 1, we assume that the subject simply compares the maximum EU over these 6 virtual lotteries with the utility of the deterministic offer in round 1. To calculate EU and utility of the offer one needs to know the parameters of the utility function, but these are just 6 EU evaluations and 1 utility evaluation. These evaluations can be undertaken within a likelihood function evaluator, given candidate values of the parameters of the utility function.

The same process can be repeated in round 2, generating another set of 5 virtual lotteries to be compared to the actual bank offer in round 2. This simulation would not involve opening as many boxes, but the logic is the same. Similarly for rounds 3 through 6. Thus for each of round 1 through 6, we can compare the utility of the actual bank offer with the maximum EU of the virtual lotteries for that round, which in turn reflects the EU of receiving a bank offer in future rounds in the underlying game.

In addition, there exists a virtual lottery in which the subject says "No Deal" in every round. This is the virtual lottery that we view as being realized in round 7 in Table 1.

There are several advantages of this approach. First, we can directly see that the contestant that has a short horizon behaves in essentially the same manner as the contestant that has a longer horizon, and just substitutes different virtual lotteries into their latent EUT calculus. This makes it easy to test hypotheses about the horizon that contestants use, although here we assume that contestants evaluate the full horizon of options available. Second, one can specify mixture models of different horizons, and let the data determine what fraction of the sample employs which horizon. Third, the approach generalizes for any known offer function, not just the ones assumed here and in Table 1. Thus it is not as specific to the *DOND* task as it might initially appear. This is important if one views *DOND* as a canonical task for examining fundamental methodological aspects of dynamic choice behavior. Those methods should not exploit the specific structure of *DOND*, unless there is no loss in generality.[14] In fact, other versions of *DOND* can be used to illustrate the flexibility of this

[14] For example, although the computational cost of taking all 6 virtual lotteries into the likelihood function for evaluation is trivial, one might be tempted to just take 2 or 3 in. In *DOND* there is a monotonicity in the bank offer fraction as each round goes by, such that little would be lost by simply using the myopic virtual lottery (the virtual lottery looking ahead just one round) and the virtual lottery looking ahead the maximal number of rounds. In general, where the option value is not known to be monotone, one cannot use such numerical short-cuts.

approach, since they sometimes employ "follow on" games that can simply be folded into the virtual lottery simulation.[15] Finally, and not least, this approach imposes virtually no numerical burden on the maximum likelihood optimization part of the numerical estimation stage: all that the likelihood function evaluator sees in a given round is a non-stochastic bank offer, a handful of (virtual) lotteries to compare it to given certain proposed parameter values for the latent choice model, and the actual decision of the contestant to accept the offer or not. This parsimony makes it easy to examine alternative specifications of the latent dynamic choice process, as illustrated below.

All estimates allow for the possibility of correlation between responses by the same subject, so the standard errors on estimates are corrected for the possibility that the responses are clustered for the same subject. The use of clustering to allow for "panel effects" from unobserved individual effects is common in the statistical survey literature.[16]

C. Estimates

The estimates from the EP function (1) are shown in Panel A of Table 2, along with estimates of the implied RRA for income levels between £10 and £250,000.[17] Since $\alpha>0$ and is statistically significantly different from zero at a *p*-value of less than 0.001, we reject CRRA.[18] On the

---

[15] For example, Mulino, Scheelings, Brooks and Faff [2006] and De Roos and Sarafidis [2006] discuss the inferential importance of the *Chance* and *Supercase* variants in the Australian version.

[16] Clustering commonly arises in national field surveys from the fact that physically proximate households are often sampled to save time and money, but it can also arise from more homely sampling procedures. For example, Williams [2000; p.645] notes that it could arise from dental studies that "collect data on each tooth surface for each of several teeth from a set of patients" or "repeated measurements or recurrent events observed on the same person." The procedures for allowing for clustering allow heteroskedasticity between and within clusters, as well as autocorrelation within clusters. They are closely related to the "generalized estimating equations" approach to panel estimation in epidemiology (see Liang and Zeger [1986]), and generalize the "robust standard errors" approach popular in econometrics (see Rogers [1993]). Wooldridge [2003] reviews some issues in the use of clustering for panel effects, noting that significant inferential problems may arise with small numbers of panels. In the *DOND* literature, De Roos and Sarafidis [2006] demonstrate that alternative ways of correcting for unobserved individual heterogeneity (random effects or random coefficients) generally provide similar estimates, but that they are quite different from estimates that ignore that heterogeneity. Botti, Conte, DiCagno and D'Ippoliti [2006] also consider unobserved indiviudual heterogeneity, and show that it is statistically significant in their models (which ignore dynamic features of the game).

[17] Since RRA is non-linear function of α and *r*, it is calculated using the "delta method," which also provides estimates of standard errors (Oehlert [1992]).

[18] We find evidence of moderate risk aversion if we impose CRRA. The point estimate for RRA is estimated fairly precisely as 0.34, with a 95% confidence interval between 0.31 and 0.37.

other hand, the estimate of α is small in relation to the smallest prizes, since we normalize all prizes by the largest prize during estimation. So we infer relatively small changes in RRA between £1 and £1,000, but substantial increases as the prizes reach £100,000 and then £250,000. The utility function is estimated over lotteries that include £250,000 with a positive probability, even if no contestant has walked away with more than £120,000 so far. We are also able to reject CARA based on these estimates of the EP function, since r is estimated to be 0.124 and is significantly different from zero with a p-value less than 0.001.[19]

Thus we conclude that *the CRRA and CARA functional forms appear to be inappropriate for these data and this specification.* We find evidence of steadily *increasing* RRA in prize income over the sizeable prize domain of the show.[20]

Inspecting the underlying calculations, in every round the virtual lottery that provides the relevant point of comparison is the one that is realized in round 7. If subjects were risk neutral, or had very little risk aversion, the average payoffs shown in Table 1 make it apparent why this is the case; however, it is true even with the high levels of risk aversion we estimate for larger prizes. We can easily constrain the estimates to limit the contestant to inspecting virtual lotteries that only look ahead a certain number of rounds. Our default, of course, is to assume that the contestant looks ahead over the complete sequence of the game.

[19] Another popular flexible specification in theoretical work is the Harmonic Absolute Risk Aversion (HARA) function $u(m) = [\eta + (m/\varphi)]^{1-\varphi}$, where η and φ are parameters to be estimated. This functional form was introduced by Merton [1971], although we use the parameterization of Gollier [2001; p.26]. RRA is defined as $m/[\eta+(m/\varphi)]$, and can vary with m. This function nests CRRA when η=0 and nests CARA when $\varphi \rightarrow \infty$. One often encounters an objection that HARA only nests CARA when one parameter equals infinity. For example, Saha [1993; p.907] claims that this "...poses problems not only in econometric estimation of risk coefficients but also in risk programming or simulation analysis." But of course one can take the reciprocal of that parameter and test if it equals zero, so we see no valid objection to HARA as a flexible functional form on these grounds. If we use HARA instead of EP we also convincingly reject CRRA and CARA. HARA is not attractive in general since it does not admit risk-loving behavior.

[20] Evidence of increasing RRA generates a confound for attempts by Post, van den Assem, Baltussen and Thaler [2006] to show that prior outcomes affect risk attitudes. They claim that contestants that experience unlucky prior outcomes exhibit significantly lower levels of risk aversion. But these are the very contestants, by construction, that come to face lower prizes than other contestants. If there is increasing RRA over the domain of prizes, variables identifying those with prior losses will pick up this variation in the weighted RRA over all prizes, and hence lead one to infer a reduction in risk aversion. Thus, a "movement along" a non-constant RRA function under EP is confounded with a "shift in" RRA because of the incorrect assumption of CRRA over all prizes.

D. The Arguments of Utility

There has been some debate over whether one models utility as a function of terminal wealth (EUTw) or income (EUTi): see Cox and Sadiraj [2006] for a review. Both specifications have been popular. The EUTw specification was widely employed in the seminal papers defining risk aversion and it's application to portfolio choice. The EUTi specification has been widely employed by auction theorists and experimental economists testing EUT, and it is the specification we have employed here. We examine this issue by *estimating* the level of wealth directly as part of the maximum likelihood process, following Harrison, List and Towe [2004], Heinemann [2005] and Andersen, Harrison and Rutström [2006].[21] In the utility function (1) we replace m with Ω+m, and estimate Ω. Thus the EP function becomes

$$u(m) = [1-\exp(-\alpha(\Omega+m)^{1-r})]/\alpha \tag{1'}$$

and the likelihood function extended to include Ω.

The results are striking. The estimates of Ω in Panel B of Table 2 is £26,011, with a standard error of £13,104 and a 95% confidence interval between £327 and £51,695. This estimate is significantly greater than zero, contrary to the assumption under an EUTi specification. It is also comparable to average annual income in the United Kingdom.[22] Thus we infer that *contestants in the game show behave as if they integrate the lottery prizes with their annual income* from their everyday employment. This result is significantly different from the previous studies mentioned above that all find little or no integration with other income when "pocket money" is at stake. It is worth emphasizing that Ω was estimated, and could have been equal to zero if that was consistent with the observed choices. The estimate of Ω has a relatively wide confidence interval, consistent with the divers demographic backgrounds of contestants on the show.

---

[21] In the *DOND* literature Bombardini and Trebbi [2005], Mulino, Scheelings, Brooks and Faff [2006], Post, van den Assem, Baltussen and Thaler [2006], De Roos and Sarafidis [2006] and Deck, Jungmin and Reyes [2006] each consider this issue by assuming alternative *exogenous* income or wealth levels.

[22] Average annual final income in 2004/2005 was £22,280, and median final income was £18,210. These estimates come from Figure 2.1 of the Department for Work and Pensions, *Households Below Average Income 2004/05*, page 10 (http://www.dwp.gov.uk/asd/hbai/hbai2005/contents.asp). They report weekly figures, and we multiply these by 52.18 to get annual figures.

Allowing for an enhanced argument of the utility function affects the estimates of risk attitudes, as might be expected. The most significant effect is that *we no longer reject the CRRA specification*, since α is *not* significantly different from zero (*p*-value = 0.661). The implied levels of RRA for prizes indicates greater risk aversion for lower prizes (roughly 0.85 in panel B compared to 0.13 in panel A), and lower risk aversion for higher prizes (0.87 in panel B for the highest prize, compared to 1.73 in panel A).

E. Evaluation Horizon

We can extract from our estimation information on the horizon actually used to evaluate the alternatives to accepting the bank offer. Figure 2 displays histograms of these evaluation horizons in each round. The results are striking, and point to these contestants putting most weight on the "terminal virtual lottery" in which they say "No Deal" in every round. Thus in round 1, almost every contestant actually used the expected utility of the round 7 virtual lottery to compare to the bank offer. The expected utilities for other virtual lotteries may well have generated the same decision, but the virtual lottery for round 7 was the one actually used since it was greater than the others in terms of expected utility.

The implication is that it is *as if* subjects largely ignore the path of offers in future rounds, and only consider the effect of the boxes they open on the probabilities attached to each remaining prize. Of course, these contestants are assumed to evaluate every implicit horizon. So if these distributions had been further to the left, it would not have implied that subject had used a shorter evaluation horizon, just that they did not *need* to use a longer horizon in this instance.

This finding may vary with different versions of the game show, with different paths for the bank offer over succeeding rounds, and naturally with differences in risk attitudes. In fact, preliminary results with data from the game show in the United States, Australia and other countries points to considerable variety in the evaluation horizons actually used.

3. Alternatives to Expected Utility Theory

There are many alternatives to EUT. We explore the main bloodlines of the family tree, with an eye to identifying those variants most appropriate *a priori* for examining behavior in the *DOND* environment.

A. Rank-Dependent Preferences

One route of departure from EUT has been to allow preferences to depend on the rank of the final outcome. The idea that one could use non-linear transformations of the probabilities as a lottery when weighting outcomes, *instead* of non-linear transformations of the outcome into utility, was most sharply presented by Yaari [1987]. To illustrate the point clearly, he assumed that one employed a linear utility function, in effect ruling out any risk aversion or risk seeking from the shape of the utility function *per se*. Instead, concave (conxex) probability weighting functions would imply risk seeking (risk aversion).[23] It was possible for a given decision maker to have a probability weighting function with both concave and convex components, and the conventional wisdom held that it was concave for smaller probabilities and convex for larger probabilities.

The idea of rank-dependent preferences had two important precursors.[24] In economics Quiggin [1982][1993] had formally presented the general case in which one allowed for subjective probability weighting in a rank-dependent manner and allowed non-linear utility functions. This branch of the family tree of choice models has become known as Rank-Dependent Utility (RDU). The Yaari [1987] model can be seen as a pedagogically important special case, and can be called Rank-Dependent Expected Value (RDEV). The other precursor, in psychology, is Lopes [1984]. Her concern was motivated by clear preferences that experimental subjects exhibited for lotteries with

[23] Camerer [2005; p.130] provides a useful reminder that "Any economics teacher who uses the St. Petersburg paradox as a "proof" that utility is concave (and gives students a low grade for not agreeing) is confusing the sufficiency of an explanation for its necessity."

[24] Of course, many others recognized the basic point that the distribution of outcomes mattered for choice in some holistic sense. Allais [1979; p.54] was quite clear about this, in a translation of his original 1952 article in French. Similarly, in psychology it is easy to find citations to kindred work in the 1960's and 1970's by Lichtenstein, Coombs and Payne, inter alia.

the same expected value but alternative shapes of probabilities, as well as the verbal protocols those subjects provided as a possible indicator of their latent decision processes. One of the most striking characteristics of *DOND* is that it offers contestants a "long shot," in the sense that there are small probabilities of extremely high prizes, but higher probabilities of lower prizes.[25]

Formally, to calculate decision weights under RDU one replaces expected utility

$$EU_i = \sum_{k=1, 20} [\, p_k \times u_k \,]. \tag{2}$$

with RDU

$$RDU_i = \sum_{k=1, 20} [\, w_k \times u_k \,]. \tag{2'}$$

where

$$w_i = \omega(p_i + ... + p_n) - \omega(p_{i+1} + ... + p_n) \tag{6a}$$

for i=1,... , n-1, and

$$w_i = \omega(p_i) \tag{6b}$$

for i=n, where the subscript indicates outcomes ranked from worst to best, and where $\omega(p)$ is some probability weighting function.

Picking the right probability weighting function is obviously important for RDU specifications. A weighting function proposed by Tversky and Kahneman [1992] has been widely used.[26] It is assumed to have well-behaved endpoints such that $\omega(0)=0$ and $\omega(1)=1$ and to imply weights

$$\omega(p) = p^\gamma / [\, p^\gamma + (1-p)^\gamma \,]^{1/\gamma} \tag{7}$$

for $0<p<1$. The normal assumption, backed by a substantial amount of evidence reviewed by Gonzalez and Wu [1999], is that $0<\gamma<1$. This gives the weighting function an "inverse S-shape,"

---

[25] Andersen, Harrison, Lau and Rutström [2006] examine the application of the resulting psychological model to these *DOND* data.

[26] Bombardini and Trebbi [2005] and De Roos and Sarafidis [2006] unfortunately use the power probability weighting function $\omega(p) = p^\varsigma$, which imposes the restriction that weighted probabilities be (weakly) concave or (weakly) convex functions *for all p*, but not both in different regions of probability space. Parametric parsimony cannot be the reason for using this function, since the functional form (7) only uses one parameter and allows concave and convex regions. We note some substantive implications of this restriction later. Blavatskyy and Pogrebna [2006c] and Botti, Conte, DiCagno and D'Ippoliti [2006] employ the same functional form (7) that we use.

characterized by a concave section signifying the overweighting of small probabilities up to a crossover-point where $\omega(p)=p$, beyond which there is then a convex section signifying underweighting. Under the RDU assumption about how these *probability* weights get converted into *decision* weights, $\gamma<1$ implies *over*weighting of extreme outcomes. Thus the probability associated with an outcome does not directly inform one about the decision weight of that outcome. If $\gamma>1$ the function takes the less conventional "S-shape," with convexity for smaller probabilities and concavity for larger probabilities.[27] Under RDU $\gamma>1$ implies *under*weighting of extreme outcomes.

We assume the CRRA functional form

$$u(m) = m^{1-\rho} / (1-\rho) \tag{1'}$$

for utility. The remainder of the econometric specification is the same as for the EUT model, generating

$$\nabla RDU = (RDU_{BO} - RDU_{I})/\mu \tag{3'}$$

instead of (3). The conditional log-likelihood becomes

$$\ln L^{RDU}(\rho, \gamma, \mu; y, X) = \textstyle\sum_i l_i^{RDU} = \sum_i [(\ln G(\nabla RDU) \mid y_i=1)+(\ln (1-G(\nabla RDU)) \mid y_i=0)] \tag{5'}$$

and requires the estimation of ρ, γ and μ.

For RDEV one replaces (2') with a specification that weights the prizes themselves, rather than the utility of the prizes:

$$RDEV_i = \textstyle\sum_{k=1, 20} [w_k \times m_k] \tag{2''}$$

where m_k is the k^{th} monetary prize. In effect, the RDEV specification is a special case of RDU with the constraint $\rho=0$.

Table 3 collects estimates of the RDEV and RDU models applied to *DOND* behavior. In each case we find estimates of $\gamma<1$, consistent with the usual expectations from the literature. Figure 3 displays the implied probability weighting function and decision weights. The decision weights are

[27] There are some well-known limitations of the probability weighting function (7). It does not allow independent specification of location and curvature; it has a crossover-point at $p=1/e=0.37$ for $\gamma<1$ and at $p=1-0.37=0.63$ for $\gamma>1$; and it is not increasing in p for small values of γ. Prelec [1998] and Rieger and Wang [2006] offer two-parameter probability weighting functions that exhibits more flexibility than (7), but for our purposes the standard probability weighting function is adequate.

shown for a 2-outcome lottery and then for a 5-outcome lottery, since these reflect the lotteries that a contestant faces in the last two rounds in which a bank offer is made. The rank-dependent specification assigns the greatest weight to the lowest prize, which we indicate by the number 1 even if it could be any of the original 22 prizes in the DOND domain. That is, by the time the contestant has reached the last choice round, the lowest prize might be 1 penny or it might be £100,000. In either case the RDU model assigns the greatest decision weight to it. Similarly, for K-outcome lotteries and K>2, a higher weight is given to the top prize compared to the others, although not as high a weight as for the lowest prize. Thus the two extreme outcomes receive relatively higher weight. Ordinal proximity to the extreme prizes slightly increases the weights in this case, but not by much. Again, the actual dollar prizes these decision weights apply to change with the history of each contestant.

There is evidence from the RDU estimates that the RDEV specification can be rejected, since ρ is estimated to be 0.321 and significantly greater than 0. Thus we infer that there is some evidence of concave utility, as well as probability weighting. As one might expect, constraining the utility function to be linear in the RDEV specification increases the curvature of the probability weighting function (since it results in a lower γ of 0.517, which implies a more concave and convex function than shown in Figure 3).

B. *Rank and Sign-Dependent Preferences: Cumulative Prospect Theory*

Original Prospect Theory

Kahneman and Tversky [1979] introduced the notion of sign-dependent preferences, stressing the role of the reference point when evaluating lotteries. In various forms, as we will see, Prospect Theory (PT) has become the most popular alternative to EUT. Original Prospect Theory (OPT) departs from EUT in three major ways: (a) allowance for subjective probability weighting; (b) allowance for a reference point defined over outcomes, and the use of different utility functions for gains or losses; and (c) allowance for loss aversion, the notion that the disutility of losses weighs more heavily than the utility of comparable gains.[28]

The first step is probability weighting, of the form $\omega(p)$ defined in (7). One of the central assumptions of OPT, differentiating it from later variants of PT, is that $w(p) = \omega(p)$, so that the transformed probabilities given by $\omega(p)$ are directly used to evaluate prospective utility:

$$PU_i = \sum_{k=1, 20} [\, \omega_k \times u_k \,]. \tag{2'''}$$

The second step in OPT is to define a reference point so that one can identify outcomes as gains and losses. Let the reference point be given by χ for a given subject in a given round. Consistent with the functional forms widely used in PT, we again use the CRRA functional form

$$u(m) = m^{1-\alpha} / (1-\alpha) \tag{1'}$$

when $m \geq \chi$, and

$$u(m) = -\lambda[(-m)^{1-\alpha} / (1-\alpha)] \tag{1''}$$

when $m < \chi$, and where λ is the loss aversion parameter. We use the same exponent α for the utility functions defined over gains and losses, even though the original statements of PT keep them theoretically distinct. Köbberling and Wakker [2005; §7] point out that this constraint is needed to identify the degree of loss aversion if one uses CRRA functional forms *and* does not want to make other strong assumptions (e.g., that utility is measurable only on a ratio scale).[29] Although λ is free to be less than 1 or greater than 1, most PT analysts presume that $\lambda \geq 1$, and we will impose that constraint.

The specification of the reference point is a critical parameter in PT. In many laboratory experiments it is assumed that the manner in which the task is framed to the subject defines the reference point that the subject uses. Thus, if one tells the subject that they have an endowment of $15 and that one lottery outcome is to have $8 taken from them, then the frame might be

---

[28] It should also be noted that OPT is formally reference dependent, but only as an artefact of the model being restricted to two outcomes.

[29] The estimates of the coefficient obtained by Tversky and Kahneman [1992] fortuitously happened to be the same for losses and gains, and many applications of PT assume that for convenience. The empirical methods of Tversky and Kahneman [1992] are difficult to defend, however: they report median values of the *estimates* obtained after fitting their model for each subject. The estimation for each subject is attractive if data permits, as in Hey and Orme [1992], but the *median estimate* has nothing to commend it statistically.

appropriately assumed to be \$15 and this outcome coded as a loss of \$8. But if the subject had been told, or expected, to earn only \$5 from the experimental task, would this be coded instead as a gain of \$3? The subjectivity and contextual nature of the reference point has been emphasized throughout, even though one often collapses it to the experimenter-induced frame in evaluating laboratory experiments. This imprecision in the reference point is not a criticism of PT, just a challenge to be careful assuming that it is always fixed and deterministic (see Schmidt, Starmer and Sugden [2005], Kőszegi and Rabin [2005][2006] and Andersen, Harrison and Rutström [2006]).[30]

In the case of *DOND* the task itself provides some natural reference points that evolve from round to round for each subject: the current bank offer to that subject. We use this information as our reference point, and simply pass χ to the likelihood function as data to be used to condition the evaluation of the probability of each decision. Thus the reference point χ varies with the bank offers and history of each individual subject, so we are not assuming that different subjects face the same reference point.[31]

The reference point also influences the nature of subjective probability weighting assumed, since different weights are allowed for gains and losses. Thus we again assume

$$\omega(p) = p^\gamma / [\, p^\gamma + (1-p)^\gamma \,]^{1/\gamma} \tag{7}$$

for gains, but estimate

$$\omega(p) = p^\phi / [\, p^\phi + (1-p)^\phi \,]^{1/\phi} \tag{7'}$$

---

[30] A corollary is that it might be a mistake to view loss aversion as a fixed parameter λ that does not vary with the context of the decision, *ceteris paribus* the reference point. See Novemsky and Kahneman [2005] and Camerer [2005; p.132, 133] for discussion of this concern. In the *DOND* setting, for example, it is hard to imagine any contestant exhibiting too much loss aversion in rounds 1 and 2, since very few contestants are in dire straits until round 3. The worst opening round in *DOND* history in the UK was on June 1, 2006, when one blighted contestant opened £250,000, £100,000, £75,000, £15,000 and £10,000. But even this contestant went on to win £9,800. In fact, of the three other UK contestants that have had "all red" opening rounds (12/12/2005, 1/2/2006 and 4/21/2006), only one went on to end up winning a low amount. Our modeling hands are currently full with the free parameters of PT that need to be estimated, and it would be premature to add complexities to the treatment of λ.

[31] We could also use the best offer received so far by the subject. More generally, the *DOND* environment provides a rich domain to study how reference points endogenously evolve over time for a subject, and the effects of allowing them to be stochastic. There are also many ways in which the reference point, once determined, could then be used to frame dynamic choices: see Thaler and Johnson [1990] for examples. By assuming that it is deterministic but evolves with the observable history of the subject we provide a simple specification as a baseline for future work in these directions.

for losses. It is common in empirical applications to assume $\gamma=\phi$. Given the striking asymmetry in the distribution of the monetary value of the prizes in *DOND*, we avoid such restrictions *a priori*.

The remainder of the econometric specification is the same as for the EUT model, generating

$$\nabla PU = (PU_{BO} - PU_L)/\mu \qquad (3''')$$

instead of (3). The conditional log-likelihood becomes

$$\ln L^{OPT}(\alpha, \lambda, \gamma, \phi, \mu; y, \chi) = \textstyle\sum_i [(\ln G(\nabla PU) \mid y_i=1)+(\ln (1-G(\nabla PU)) \mid y_i=0)] \qquad (5'')$$

and requires the estimation of α, λ, γ, ϕ and μ.

The primary logical problem with OPT was that it implied violations of stochastic dominance. Whenever $\gamma\neq 1$ or $\phi\neq 1$, it is possible to find non-degenerate lotteries such that one would stochastically dominate the other, but would be assigned a lower PU. Examples arise quickly when one recognizes that $\gamma(p_1 + p_2) \neq \gamma(p_1) + \gamma(p_2)$ for some p_1 and p_2. Kahneman and Tversky [1979] dealt with this problem by assuming that evaluation using OPT only occurred after dominated lotteries were eliminated. Our model of OPT does not contain such an editing phase, but the stochastic error term μ could be interpreted as a reduced form proxy for that editing process.[32]

Cumulative Prospect Theory

The notion of rank-dependant decision weights was incorporated into OPT by Starmer and Sugden [1989], Luce and Fishburn [1991] and Tversky and Kahneman [1992]. Instead of implicitly assuming that $w(p) = \omega(p)$, it allowed $w(p)$ to be defined as in the RDU specification (6a) and (6b). The sign-dependence of subjective probability weighting in OPT, leading to the estimation of

[32] In other words, evaluating the PU of two lotteries, without having edited out dominated lotteries, might lead to a dominated lottery having a higher PU. But if subjects always reject dominated lotteries, the choice would appear to be an error to the likelihood function. Apart from searching for better parameters to explain this error, as the maximum likelihood algorithm does as it tries to find parameter estimates that reduce any other prediction error, our specification allows μ to increase. We stress that this argument is not intended to rationalize the use of separable probability weights in OPT, just to explain how a structural model with stochastic errors might account for the effects of stochastic dominance. Wakker [1989] contains a careful account of the notion of transforming probabilities in a "natural way" but without violating stochastic dominance.

different probability weighting functions (7) and (7') for gains and losses, is maintained in Cumulative Prospect Theory (CPT). Thus there is a separate decumulative function used for gains and losses, but otherwise the logic is the same as for RDU.

Table 4 reports estimates from the CPT specifications, and Figure 4 displays the estimated probability weighting functions and decision weights. Contestants exhibit a significantly concave (convex) utility function over gains (losses), with α=0.47.

There is no evidence of loss aversion, with λ estimated to be 1.[33] Since this result might be viewed as heretical by some, it is important to be clear how it is obtained. As noted earlier, we constrain $\lambda \geq 1$. Such constraints are implemented by estimating a surrogate parameter τ that is unconstrained, and scaling it with fixed parameters A and B to actually take on the value A + (B-A) × [1/(1+exp(τ))]. Thus, as τ gets arbitrarily large the original parameter λ converges to A, and as τ gets arbitrarily small the original parameter λ converges to B. In this instance we set A=1 and B=8, although the upper bound is immaterial here. Our CRRA specification deliberately followed the recommendations of Köbberling and Wakker [2005; §7] in order that one could define a theoretically coherent and general index of loss aversion. However, one might suspect that the CRRA functional form has some idiosyncrasy that affects this conclusion. We therefore followed their recommendation to use a CARA functional form instead, and obtained exactly the same result.

The pattern of subjective probability weighting for losses and gains is shown in Figure 4.

---

[33] Blavatskyy and Pogrebna [2006a] claim that there is no evidence of loss aversion in the Italian version of the *DOND* game show. They examine the frequency with which contestants agree to swap "their box" for another unopened box when given the opportunity. They argue that a loss averse contestant would never agree to such a swap since the distribution of prizes remains the same, and a swap opens the subject up to some losses. Implicit in this claim is the assumption that the current distribution of (unopened) prizes is the reference point that the subject uses *for this binary decision*, which may or may not be true. Certainly the reference point for this "swap" decision need not be the same as the reference point for the "Deal" decision. But the presumption that the current endowment is the appropriate reference point for comparable "swap" choices in lab experiments is commonly made when testing the so-called "endowment effect" (e.g., Plott and Zeiler [2005]). In any event, 61% of subjects in 100 shows rejected the first offer to swap, which is required as part of the rules of the Italian version. If one accepts the premiss of the test for loss aversion, this is still evidence that 61% of the contestants acted consistently with loss aversion. Of course, under their premiss, EUT contestants are indifferent, so 61% compared to 50% with this size sample is hardly convincing evidence one way or the other. Our test is more direct by estimating the structural parameter λ actually defining loss aversion jointly with the other parameters of the model. On the other hand, our test does rely on the set of parametric assumptions made about functional forms.

The pattern of probability weighting is standard for gains, since $\gamma<1$ resulting in overweighting of small probabilities and underweighting of probabilities greater than ⅓. The pattern is non-standard for losses, with $\phi>1$, but there is relatively little probability weighting over losses under CPT. In *DOND* it is only in the final round 6 that a subject would face probabilities on outcomes greater than 0.2 so the predominant pattern appears to be overweighting.[34] Figure 3 also shows the effects of probability weighting on decision weights under CPT. For gains we see the same pattern as obtained with the RDU specification: each extreme outcome receives greater weight, and the "interior" outcomes receive roughly the same weights. This latter result is striking as one evaluates decision weights for lotteries with more and more outcomes. For losses there is virtually no effect from probability weighting, with roughly the same probability weighting that one would obtain under EUT.

In general, the *DOND* environment reminds us that one must take a structural perspective when estimating CPT models. Estimates of the loss aversion parameter depend intimately on the assumed reference point, as one would expect since the latter determines what are to be viewed as losses. So if we have assumed the wrong reference point, we will not reliably estimate the degree of loss aversion. However, we do not get loss aversion leaping out at us when we make a natural assumption about the reference point, and this points to a key operational weakness of CPT: the need to specify what the reference point is. Loss aversion *may* be present for *some* reference point, but it is not present for the one we used, and none others are "obviously" better. Without a convincing argument about the correct reference point, and evidence for loss aversion conditional on that reference point, one simply cannot claim that loss aversion is always present. Thus our

---

[34] Bombardini and Trebbi [2005] and De Roos and Sarafidis [2006] use the power probability weighting function $\omega(p) = p^{\varsigma}$, as noted earlier. For values of the exponent $\varsigma<1$ ($\varsigma>1$), this implies overweighting (underweighting) *for all p*. Hence, if probabilities for prizes in the last three rounds, the sample used by Bombardini and Trebbi [2005], are between ¼ and ½, they would be forcing the power function to fit "data that wants to be concave and then convex." It is therefore not a surprise that they estimate $\varsigma=0.92$ with a standard error of 0.06, and *cannot* reject the EUT-consistent null that $\varsigma=1$. This is an artefact of assuming the restrictive power function, not an inference solely from the data. In addition, they each estimate RDU models that assume that all decision weights are in the gain domain. If there is, in fact, a combination of gain domain and loss domain choices then their estimates would be a weighted average of the distinct functions we estimate, further constraining the probability weighting in the gain domain to appear to be less concave over small probabilities than it actually is.

results point to a significant general problem making CPT operational in field settings. This specification ambiguity is arguably not present in the lab, where one can frame tasks, but is a serious issue in the field. Similarly, estimates of the nature of probability weighting vary with changes in reference points, loss aversion parameters, and the concavity of the utility function, and *vice versa*. All of this is to be expected from the CPT model, but necessitates joint econometric estimation of these parameters if one is to be able to make consistent statements about behavior.[35]

4. Additional Research Questions

We believe that the task environment in the *DOND* game show raises many exciting questions for the formal modeling of dynamic choice behavior. We consider three open issues.

The first open issue is the interpretation of the reference point in a dynamic setting. Is the reference point solely driven by the induced framing of the task to the subject by the game show rules or experimenter, or is there is some subjective, endogenous, history-dependent, "creative" process being used to define it (Thaler and Johnson [1990])? Similarly, is the reference point best characterized as a deterministic point of reference, or as a lottery itself (Sugden [2003], Schmidt, Starmer and Sugden [2005], and Kőszegi, and Rabin [2005][2006])? These questions are related, of course. If there is a role for subjective expectations in forming the reference point (e.g., "how much do people normally get offered in round 3 of this game?"), then it is hard to see how one can assume

[35] In addition, it is essential to specify a proper likelihood function. Post, van den Assem, Baltussen and Thaler [2006] use a grid-search over all of the parameters of a CPT specification that is similar to ours, which allows them to estimate these parameters jointly. However, they choose best estimates based on the number of "hits" they get with a certain set of parameters, where a hit is simply a correct prediction. But a correct prediction is not the same thing as the *probability* of a correct prediction, which is what the likelihood is for discrete data such as "Deal" or "No Deal" choices. Thus one cannot say if a "hit" involves probabilities of the correct prediction that are 0.51 or 0.99. It is therefore entirely possible that a parameter set with a lower average hit rate could have a much higher likelihood. Similarly, there is no coherent statistical basis for saying if a hit rate of 69% (CPT with their preferred estimates) is better than 65% (CPT assuming parameter values from Tversky and Kahneman [1992]) or 61% (their EUT specification). However, they do claim that the hit rate of 69% "... is substantially higher than the 60-61% for the expected utility model" (p.28). In fact, the parameter estimate they use for their homogenous EUT model is not chosen with a grid search to maximize the hit rate, but is set equal to the average from a heterogeneous EUT specification. That heterogeneous EUT specification itself has some serious problems, discussed in an appendix (available on request). One certainly cannot claim that apples are being compared with apples when the 61% from EUT is compared to the 69% from CPT, or that any such comparison makes statistical sense.

that the reference point is deterministic.

The second open issue is the role of "editing" in the choice process. Editing plays an important role in most psychologically-based models of the decision-making process, but has been completely dropped in the modern versions of CPT that are so popular. The famous "jump discontinuity" of the probability weighting function in Kahneman and Tversky [1979; Figure 4, p.283] no doubt upset the smooth sensibilities of those wanting tidy mathematical formulae, but was motivated by a long tradition of the study of "task representation" in cognitive psychology.[36] In particular, work on similarity relations would seem to fit in perfectly to the *DOND* environment (e.g., Tversky [1977], Luce [1956], Rubinstein [1988] and Leland [1994]). Consider the prizes in the UK version of *DOND*. The smallest four prizes are 1p, 10p, 50p and £1, which are arguably the same as zero to the contestant. The largest four prizes are £50,000, £75,000, £100,000 and £250,000. These are not at all obviously similar in the same sense that the bottom prizes are. With due regard for the effects of inflation on his numerical example, Luce [1956; p.190] makes a related point with respect to experimental tests of his own model of imperfectly discriminable utility:

> It seems extremely doubtful that laboratory experiments of the type so far reported in the literature, which for the most part have involved simple gambling situations, can utilize our model. These designs have been carefully cultivated and abstracted from actual social situations in such a manner that the preference relation is almost certainly a weak order because it reflects a natural ordering of money by amounts. It is hard to imagine anyone in such a laboratory experiment who will be indifferent between $1 and [a lottery offering $1 with probability α and $0 with probability (1-$\alpha$)] for any $\alpha>0$. One senses, however, that in society there are many decision situations with which we wish to deal in which the indifference relation is more complex than an equivalence relation. If this intuition is correct, it should be a challenge to the experimenter to try to reproduce these situations as simple laboratory experiments where the nature and origin of the intransitivities can be carefully examined.

Our point is that the distribution of prizes in the naturally occurring *DOND* game show provide such an experiment, with the advantage that the prizes are all money. That distribution has two features: a vast range, and a marked skewness (the average prize is £25,712 but the median is only

[36] Camerer [2005; p.130] calls it "... probably the only deliberately discontinuous function published in *Econometrica* – and with good reason."

£875).[37] A formal analysis that examines editing processes in this setting would be valuable, and the naturally occurring *DOND* environment encourages that analysis.

The third open issue is the manner in which one characterizes heterogeneity of preferences. The natural experiment does not provide a long list of observable characteristics for each contestant. This requires that one either pool over subjects under the assumption that they have the same preferences, as we have done; make restrictive assumptions that allow one to identify bounds for a given contestant, but then provide contestant-specific estimates (e.g., Post, van den Assem, Baltussen and Thaler [2006]); or pay more attention to statistical methods that allow for unobserved heterogeneity. One such method is to allow for random coefficients of each structural model to represent an underlying variation in preferences across the sample (e.g., Train [2003; ch.6], De Roos and Sarafidis [2006] and Botti, Conte, DiCagno and D'Ippoliti [2006]). This is quite different from allowing for standard errors in the pooled coefficient, as we have done. Another method is to allow for finite mixtures of alternative structural models, recognizing that some choices or subjects may be better characterized in this domain by one latent decision-making process and that others may be better characterized by some other process (e.g., Harrison and Rutström [2005]). These methods are not necessarily alternatives, but they each demand relatively large data sets and considerable attention to statistical detail.

5. Conclusions

The *Deal Or No Deal* paradigm is important for several reasons. It incorporates many of the dynamic, forward-looking decision processes that strike one as a natural counterpart to a wide range of fundamental economic decisions in the field. The "option value" of saying "No Deal" has clear parallels to the financial literature on stock market pricing, as well as to many investment decisions that have future consequences (so-called "real options"). There is no frictionless market ready to

[37] The median here is the mid-point of the £750 and £1,000 prizes. The distribution of prizes has a coefficient of skewness is 3.06, where 0 indicates a symmetric distribution. Average *earnings*, by contrast, were £16,750 with median earnings of £13,000, which is somewhat less skewed. The coefficient of skewness for earnings is still relatively large at 2.4.

price these options, so familiar arbitrage conditions for equilibrium valuation play no immediate role, and one must worry about how the individual makes these decisions. The game show offers a natural experiment, with virtually all of the major components replicated carefully from show to show, and even from country to country.

Our analysis, as well as complementary studies of different *Deal Or No Deal* data, show the importance of valuation heterogeneity across individuals, as well as the sensitivity to the correct stochastic modeling of the continuation value of current choices. We confirm many of the results from careful laboratory experiments of Holt and Laury [2002][2005] using EUT under varying stake conditions: one must account for varying RRA to explain behavior. However, we also show that the arguments of the utility function in this large-stake domain are not just the prizes of the lotteries, and that CRRA becomes a reasonable characterization if one allows for a more flexible specification of those arguments.

This environment should be a natural breeding ground for many of the informal and formal models that have developed to consider the role of history and framing on choices under uncertainty. Our analysis shows that there is a major role for several alternatives to EUT, in particular the role accorded probability weighting. But we do not find evidence of loss aversion, using a full maximum likelihood estimation of a model of CPT-drive behavior. This result is conditional on a specification of the reference point which strikes us as natural, and there may exist other reference points that generate loss aversion (although we have not found any). Without more guidance from the theory as to what reference point is correct it is hard to undertake reliable tests of the properties of that theory in such naturally occurring domains. The *DOND* task domain provides a fertile terrain for examining extensions to the models considered here.

Table 1: Virtual Lotteries for British *Deal or No Deal* Game Show

Average (standard deviation) of payoff from virtual lottery using 10,000 simulations

Round	Active Contestants	Deal!	Average Offer	Looking at virtual lottery realized in ... round 2	round 3	round 4	round 5	round 6	round 7
1	211 100%	0	**£4,253**	**£7,422** (£7,026)	**£9,704** (£9,141)	**£11,528** (£11,446)	**£14,522** (£17,118)	**£19,168** (£32,036)	**£26,234** (£52,640)
2	211 100%	0	**£6,909**		**£9,705** (£9,123)	**£11,506** (£11,409)	**£14,559** (£17,206)	**£19,069** (£31,935)	**£26,118** (£52,591)
3	211 100%	8	**£9,426**			**£11,614** (£11,511)	**£14,684** (£17,379)	**£19,304** (£32,533)	**£26,375** (£53,343)
4	203 96%	45	**£11,836**				**£15,457** (£18,073)	**£20,423** (£34,230)	**£27,892** (£55,678)
5	158 75%	75	**£13,857**					**£18,603** (£31,980)	**£25,282** (£51,845)
6	83 39%	38	**£13,861**						**£18,421** (£41,314)
7	45 21%								

Note: Data drawn from observations of 211 contestants on the British game show, plus author's simulations of virtual lotteries as explained in the text.

Figure 1: Two Virtual Lottery Distributions in Round 1

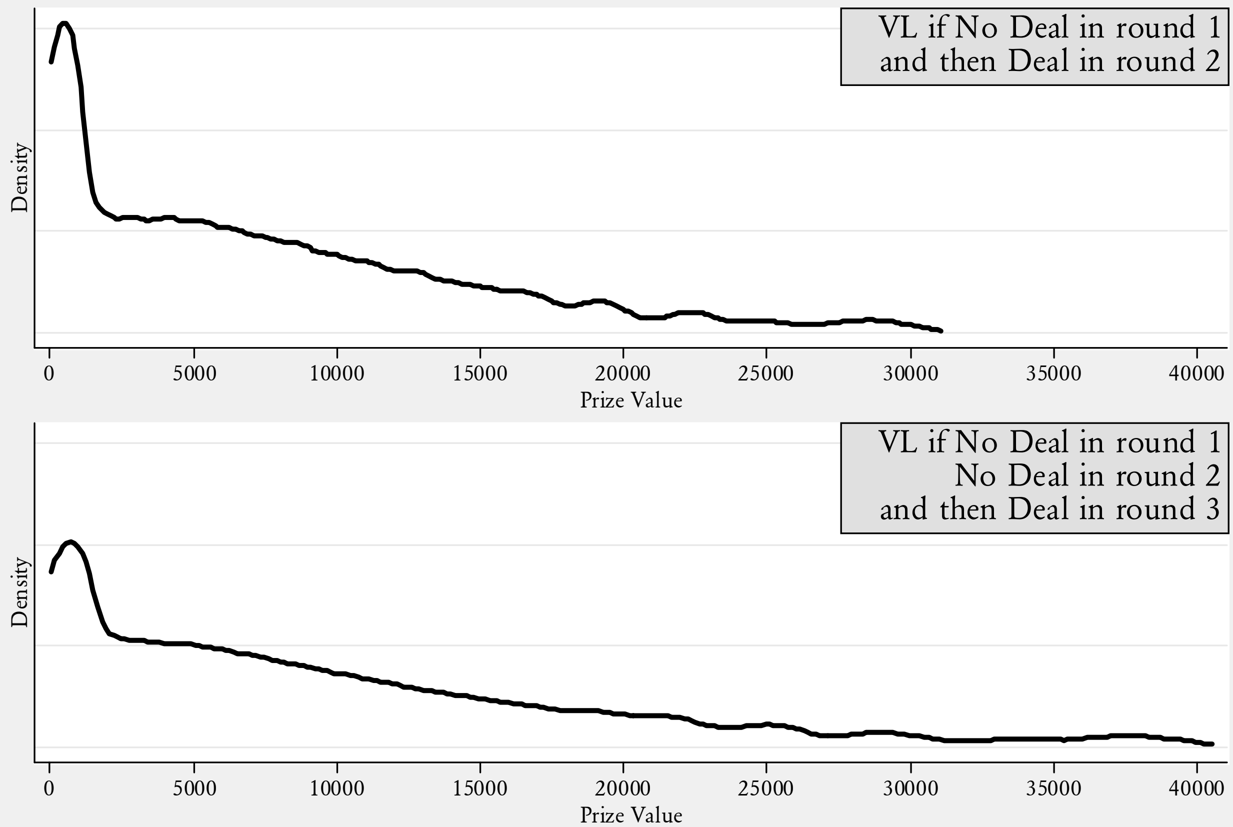

Table 2: Estimates Assuming EUT

Parameter	Estimate	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
A. EUT, assuming utility only defined over prizes: $u(m) = [1-\exp(-\alpha m^{1-r})]/\alpha$				
r	0.124	0.0005	0.123	0.125
α	1.832	0.0010	1.830	1.834
μ	0.024	0.001	0.021	0.026
		Relative Risk Aversion		
£1	0.124	0.0005	0.123	0.125
£10	0.125	0.0005	0.123	0.126
£100	0.126	0.0005	0.125	0.127
£1,000	0.137	0.0006	0.136	0.138
£10,000	0.220	0.0007	0.219	0.221
£100,000	0.844	0.0008	0.842	0.845
£250,000	1.729	0.0005	1.727	1.729
B. EUT, assuming utility defined over prizes and income: $u(m) = [1-\exp(-\alpha(\Omega+m)^{1-r})]/\alpha$				
r	0.843	0.192	0.466	1.219
α	0.147	0.335	-0.510	0.803
Ω	£26,011	£13,104	£327	£51,695
μ	0.014	0.019	-0.025	0.052
		Relative Risk Aversion ‡		
£1	0.846	0.203	0.448	1.244
£10	0.847	0.206	0.443	1.251
£100	0.849	0.209	0.439	1.259
£1,000	0.852	0.212	0.436	1.269
£10,000	0.857	0.215	0.435	1.278
£100,000	0.863	0.216	0.438	1.287
£250,000	0.866	0.216	0.442	1.290

‡ For comparability these estimates of RRA assume Ω=£0. If one instead assumes Ω=£26,011, the values range from 0.86 to 0.87, with 95% confidence intervals of 0.43 and 1.28.

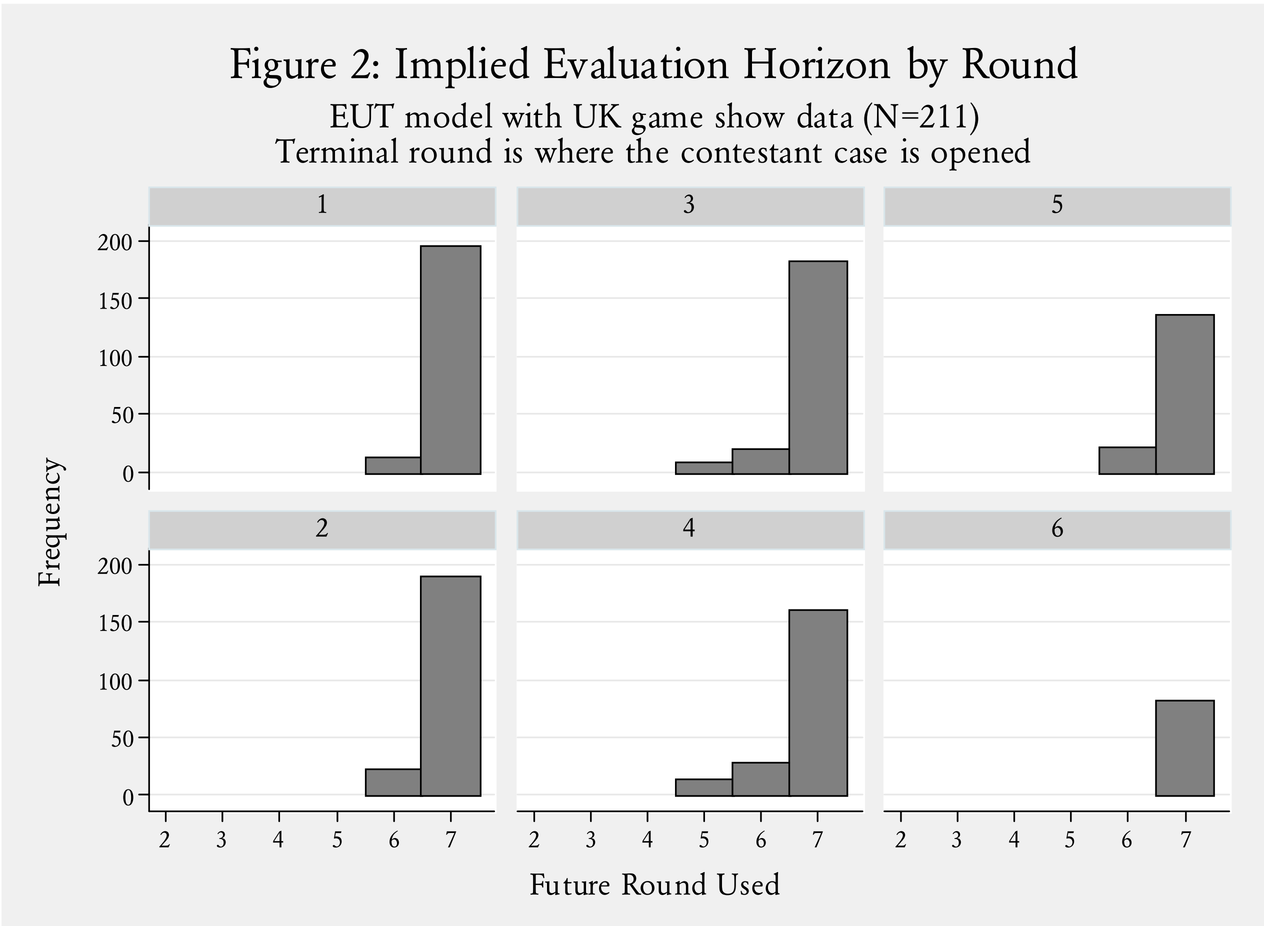
Figure 2: Implied Evaluation Horizon by Round
EUT model with UK game show data (N=211)
Terminal round is where the contestant case is opened
1
3
5
2
4
6
200
150
100
50
0
Frequency
2 3 4 5 6 7
Future Round Used

Table 3: Estimates Assuming RDU

Parameter	Estimate	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
A. RDEV, assuming utility is linear in prizes				
γ	0.517	0.133	0.256	0.779
μ	0.149	0.036	0.078	0.221
B. RDU				
ρ	0.321	0.038	0.246	0.396
γ	0.667	0.052	0.565	0.769
μ	0.256	0.033	0.191	0.321

Figure 3: Decision Weights under RDU

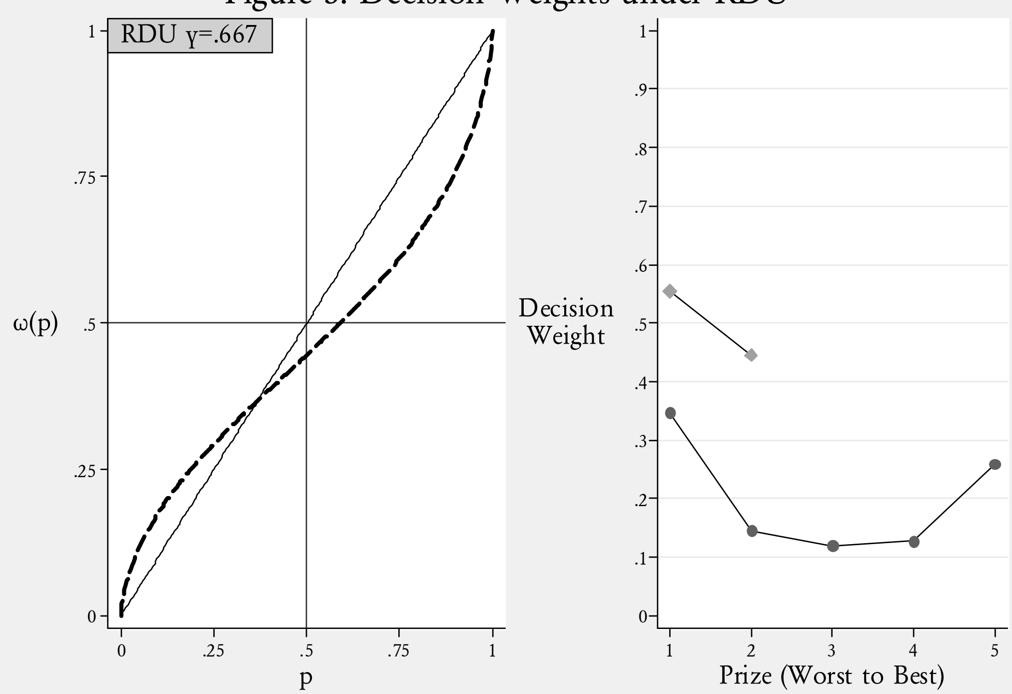

Table 4: Estimates Assuming CPT

Parameter	Estimate	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
α	0.467	0.032	0.405	0.529
λ	1.000	†		
γ	0.635	0.054	0.530	0.741
ϕ	1.129	0.241	0.657	1.602
μ	0.338	0.030	0.280	0.396

† Estimates are at lower bound of a constraint, so standard errors cannot be reliably estimated.

Figure 4: Decision Weights under CPT

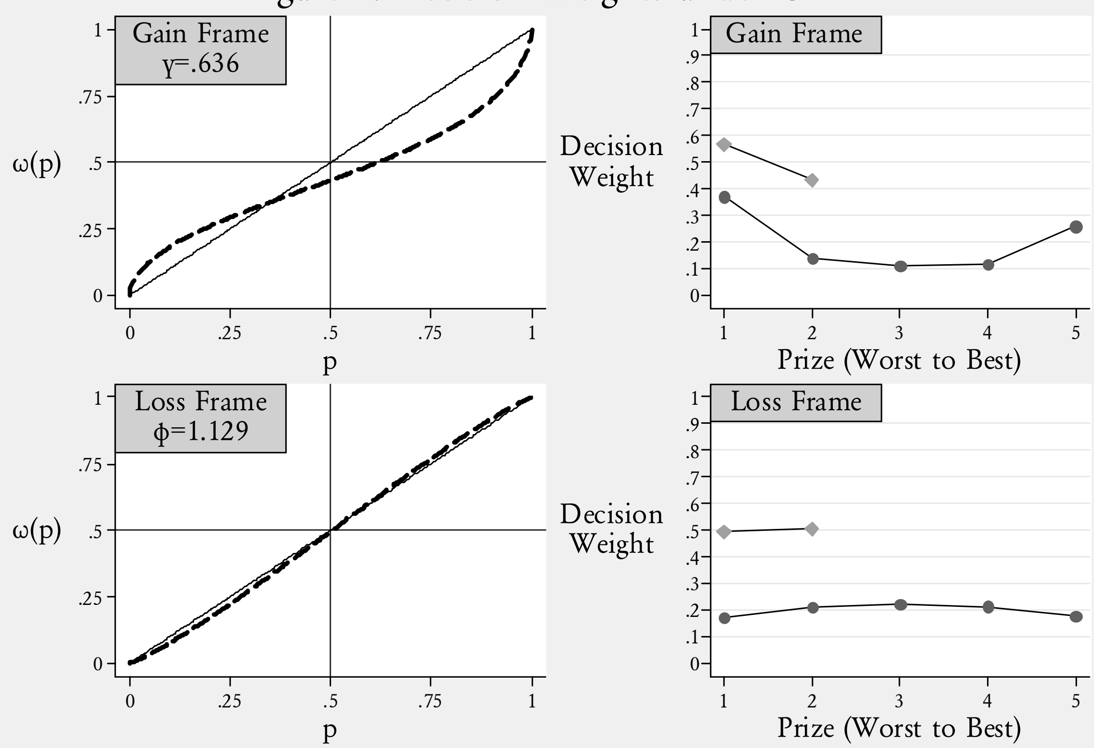

References

Abdellaoui, Mohammed; Barrios, Carolina, and Wakker, Peter P., "Reconciling Introspective Utility With Revealed Preference: Experimental Arguments Based on Prospect Theory," Department of Economics, University of Amsterdam, April 2005; forthcoming, *Journal of Econometrics*.

Allais, Maurice, "The Foundations of Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School," in M. Allais & O. Hagen (eds.), *Expected Utility Hypotheses and the Allais Paradox* (Dordrecht, the Netherlands: Reidel, 1979).

Andersen, Steffen; Harrison, Glenn W., Lau, Morten I., and Rutström, E. Elisabet, "Dual Criteria Decisions," *Working Paper 06-11*, Department of Economics, College of Business Administration, University of Central Florida, 2006.

Andersen, Steffen; Harrison, Glenn W., and Rutström, E. Elisabet, "Choice Behavior, Asset Integration and Natural Reference Points," *Working Paper 06-04*, Department of Economics, College of Business Administration, University of Central Florida, 2006.

Ballinger, T. Parker, and Wilcox, Nathaniel T., "Decisions, Error and Heterogeneity," *Economic Journal*, 107, July 1997, 1090-1105.

Baltussen, Guido; Post, Thierry, and van den Assem, Martijn, "Stakes, Prior Outcomes and Distress in Risky Choice: An Experimental Study Based on `Deal or No Deal'," *Working Paper*, Department of Finance, Erasmus School of Economics, Erasmus University, August 2006.

Beetsma, R.M.W.J., and Schotman, P.C., "Measuring Risk Attitudes in a Natural Experiment: Data from the Television Game Show *Lingo*," *Economic Journal*, 111(474), 2001, 821-848.

Berk, Jonathan B.; Hughson, Eric, and Vandezande, Kirk, "The Price Is Right, But Are the Bids? An Investigation of Rational Decision Theory," *American Economic Review*, 86(4), September 1996, 954-970.

Blavatskyy, Pavlo, and Pogrebna, Ganna, "Loss Aversion? Not with Half-a-Million on the Table!" *Working Paper 274*, Institute for Empirical Research in Economics, University of Zurich, July 2006a.

Blavatskyy, Pavlo, and Pogrebna, Ganna, "Testing the Predictions of Decision Theories in a Natural Experiment When Half a Million Is At Stake," *Working Paper 291*, Institute for Empirical Research in Economics, University of Zurich, June 2006b.

Blavatskyy, Pavlo, and Pogrebna, Ganna, "Testing Models of Stochastic Choice in *Deal or No Deal*," *Working Paper*, Institute for Empirical Research in Economics, University of Zurich, September 2006c.

Bombardini, Matilde, and Trebbi, Francesco, "Risk Aversion and Expected Utility Theory: A Field Experiment with Large and Small Stakes," *Working Paper 05-20*, Department of Economics, University of British Columbia, November 2005.

Botti, Fabrizio; Conte, Anna; DiCagno, Daniela, and D'Ippoliti, Carlo, "Risk Attitude in Real Decision Problems," *Unpublished Manuscript*, LUISS Guido Carli, Rome, September 2006.

Camerer, Colin, "Three Cheers – Psychological, Theoretical, Empirical – for Loss Aversion," *Journal* *of Marketing Research*, 42, May 2005, 129-133.

Campbell, John Y.; Lo, Andrew W., and MacKinlay, A. Craig, *The Econometrics of Financial Markets* (Princeton: Princeton University Press, 1997).

Cox, James C., and Sadiraj, Vjollca, "Small- and Large-Stakes Risk Aversion: Implications of Concavity Calibration for Decision Theory," *Games & Economic Behavior*, 56(1), July 2006, 45-60.

Cubitt, Robin P., and Sugden, Robert, "Dynamic Decision-Making Under Uncertainty: An Experimental investigation of Choices between Accumulator Gambles," *Journal of Risk & Uncertainty*, 22(2), 2001, 103-128.

Deck, Cary; Lee, Jungmin, and Reyes, Javier, "Risk Attitudes in Large Stake Gambles: Evidence from a Game Show," *Working Paper*, Department of Economics, University of Arkansas, 2006.

De Roos, Nicolas, and Sarafidis, Yianis, "Decision Making Under Risk in *Deal or No Deal*," *Working Paper*, School of Economics and Political Science, University of Sydney, April 2006.

Février, Philippe, and Linnemer, Laurent, "Equilibrium Selection: Payoff or Risk Dominance? The Case of the 'Weakest Link'," *Journal of Economic Behavior & Organization*, 60(2), June 2006, 164-181.

Gertner, R., "Game Shows and Economic Behavior: Risk-Taking on *Card Sharks*," *Quarterly Journal of Economics*, 108(2), 1993, 507-521.

Gollier, Christian, *The Economics of Risk and Time* (Cambridge, MA: MIT Press, 2001).

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

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

Harrison, Glenn W.; List, John A., and Towe, Charles, "Naturally Occurring Preferences and Exogenous Laboratory Experiments: A Case Study of Risk Aversion," *Working Paper 04-22*, Department of Economics, College of Business Administration, University of Central Florida, 2004; forthcoming, *Econometrica*.

Harrison, Glenn W., and Rutström, E. Elisabet, "Expected Utility Theory *and* Prospect Theory: One Wedding and A Decent Funeral," *Working Paper 05-18*, Department of Economics, College of Business Administration, University of Central Florida, 2005.

Hartley, Roger; Lanot, Gauthier, and Walker, Ian, "Who Really Wants to be a Millionaire? Estimates of Risk Aversion from Gameshow Data," *Working Paper*, Department of Economics, University of Warwick, February 2005.

Healy, Paul, and Noussair, Charles, "Bidding Behavior in the *Price Is Right* Game: An Experimental Study," *Journal of Economic Behavior & Organization*, 54, 2004, 231-247.

Heinemann, Frank, "Measuring Risk Aversion and the Wealth Effect," *Working Paper*, University of Munich, September 2005.

Hey, John, "Experimental Investigations of Errors in Decision Making Under Risk," *European Economic Review*, 39, 1995, 633-640.

Hey, John D., "Experimental Economics and the Theory of Decision Making Under Uncertainty," *Geneva Papers on Risk and Insurance Theory*, 27(1), June 2002, 5-21.

Hey, John D., and Orme, Chris, "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62(6), November 1994, 1291-1326.

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

Kahneman, Daniel, and Tversky, Amos, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 1979, 263-291.

Köbberling, Veronika, and Wakker, Peter P., "An Index of Loss Aversion," *Journal of Economic Theory*, 122, 2005, 119-131.

Kőszegi, Botond, and Rabin, Matthew, "Reference-Dependent Risk Attitudes," *Unpublished Manuscript*, Department of Economics, University of California at Berkeley, October 2005.

Kőszegi, Botond, and Rabin, Matthew, "A Model of Reference-Dependent Preferences," *Quarterly Journal of Economics*, 121(4), November 2006, forthcoming.

Leland, W. Jonathan, "Generalized Similarity Judgements: An Alternative Explanation for Choice Anomalies," *Journal of Risk & Uncertainty*, 9, 1994, 151-172.

Levitt, Steven D., "Testing Theories of Discrimination: Evidence from the Weakest Link," *Journal of Law & Economics*, XLVII, October 2004, 431-452.

Liang, K-Y., and Zeger, S.L., "Longitudinal Data Analysis Using Generalized Linear Models," *Biometrika*, 73, 1986, 13-22.

List, John A., "*Friend or Foe*: A Natural Experiment of the Prisoner's Dilemma," *Review of Economics & Statistics*, 88(3), August 2006, forthcoming.

Loomes, Graham; Moffatt, Peter G., and Sugden, Robert, "A Microeconometric Test of Alternative Stochastic Theories of Risky Choice," *Journal of Risk and Uncertainty*, 24(2), 2002, 103-130.

Loomes, Graham, and Sugden, Robert, "Incorporating a Stochastic Element Into Decision Theories," *European Economic Review*, 39, 1995, 641-648.

Lopes, Lola L., "Risk and Distributional Inequality," *Journal of Experimental Psychology: Human Perception and Performance*, 10(4), August 1984, 465-484.

Luce, R. Duncan, "Semiorders and a Theory of Utility Discrimination," *Econometrica*, 24, 1956, 178-191.

Luce, R. Duncan, and Fishburn, Peter C., "Rank and Sign-Dependent Linear Utility Models for Finite First-Order Gambles," *Journal of Risk & Uncertainty*, 4, 1991, 29-59.

Merton, Robert C., "Optimum Consumption and Portfolio Rules in a Continuous-Time Model," *Journal of Economic Theory*, 3, 1971, 373-413.

Metrick, Andrew, "A Natural Experiment in 'Jeopardy!'," *American Economic Review*, 85(1), March 1995, 240-253.

Mulino, Daniel; Scheelings, Richard; Brooks, Robert, and Faff, Robert, "An Empirical Investigation of Risk Aversion and Framing Effects in the Australian Version of *Deal Or No Deal*," *Working Paper*, Department of Economics, Monash University, June 2006.

Novemsky, Nathan, and Kahneman, Daniel, "The Boundaries of Loss Aversion," *Journal of Marketing Research*, 42, May 2005, 119-128.

Oehlert, Gary W., "A Note on the Delta Method," *The American Statistician*, 46(1), February 1992, 27-29.

Plott, Charles R. and Zeiler, Kathryn, "The Willingness to Pay-Willingness to Accept Gap, the 'Endowment Effect,' Subject Misconceptions, and Experimental Procedures for Eliciting Valuations," *American Economic Review*, 95(3), June 2005, 530-545.

Post, Thierry; van den Assem, Martijn; Baltussen, Guido, and Thaler, Richard, "Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show," *Working Paper*, Department of Finance, Erasmus School of Economics, Erasmus University, April 2006.

Prelec, Drazen, "The Probability Weighting Function," *Econometrica*, 66, 1998, 497-527.

Quiggin, John, "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization*, 3(4), 1982, 323-343.

Quiggin, John, *Generalized Expected Utility Theory: The Rank-Dependent Model* (Norwell, MA: Kluwer Academic, 1993).

Rabin, Matthew, "Risk Aversion and Expected Utility Theory: A Calibration Theorem," *Econometrica*, 68, 2000, 1281-1292.

Rieger, Marc Oliver and Wang, Mei, "Cumulative Prospect Theory and the St. Petersburg Paradox," *Economic Theory*, 28, 2006, 665-679.

Rogers, W. H., "Regression standard errors in clustered samples," *Stata Technical Bulletin*, 13, 1993, 19-23.

Rubinstein, Ariel, "Similarity and Decision-making Under Risk (Is There a Utility Theory Resolution to the Allais Paradox?)," *Journal of Economic Theory*, 46, 1988, 145-153.

Rubinstein, Ariel, "Comments on the Risk and Time Preferences in Economics," *Unpublished Manuscript*, Department of Economics, Princeton University, 2002.

Saha, Atanu, "Expo-Power Utility: A Flexible Form for Absolute and Relative Risk Aversion," *American Journal of Agricultural Economics*, 75(4), November 1993, 905-913.

Schmidt, Ulrich; Starmer, Chris, and Sugden, Robert, "Explaining Preference Reversal with Third-Generation Prospect Theory," *Working Paper*, School of Economic and Social Science, University of East Anglia, 2005.

Starmer, Chris, and Sugden, Robert, "Violations of the Independence Axiom in Common Ratio Problems: An Experimental Test of Some Competing Hypotheses," *Annals of Operational*

Research, 19, 1989, 79-102.

Sugden, Robert, "Reference-Dependent Subjective Expected Utility," *Journal of Economic Theory*, 111, 2003, 172-191.

Tenorio, Rafael, and Cason, Timothy, "To Spin or Not To Spin? Natural and Laboratory Experiments from *The Price is Right*," *Economic Journal*, 112, 2002, 170-195.

Thaler, Richard H., and Johnson, Eric J., "Gambling With The House Money and Trying to Break Even: The Effects of Prior Outcomes on Risky Choice," *Management Science*, 36(6), June 1990, 643-660.

Train, Kenneth E., *Discrete Choice Methods with Simulation* (New York: Cambridge University Press, 2003).

Tversky, Amos, "Features of Similarity," *Psychological Review*, 84(4), 1977, 327-352.

Tversky, Amos, and Kahneman, Daniel, "Advances in Prospect Theory: Cumulative Representations of Uncertainty," *Journal of Risk & Uncertainty*, 5, 1992, 297-323; references to reprint in D. Kahneman and A. Tversky (eds.), *Choices, Values, and Frames* (New York: Cambridge University Press, 2000).

Wakker, Peter P., "Transforming Probabilities Without Violating Stochastic Dominance," in E. Roskam (ed.), *Mathematical Psychology in Progress* (Berlin: Springer, 1989).

Williams, Rick L., "A Note on Robust Variance Estimation for Cluster-Correlated Data," *Biometrics*, 56, June 2000, 645-646.

Wooldridge, Jeffrey, "Cluster-Sample Methods in Applied Econometrics," *American Economic Review (Papers & Proceedings)*, 93, May 2003, 133-138.

Yaari, Menahem E., "The Dual Theory of Choice under Risk," *Econometrica*, 55(1), 1987, 95-115.

Appendix: Detailed Comments on Related Studies (NOT FOR PUBLICATION)

The previous literature has employed three types of empirical strategies with the *DOND* data.

The first is the calculation of CRRA bounds at which a given subject is indifferent between one choice or another. These bounds can be calculated for each subject and each choice, so they have the advantage of not assuming that each subject has the same risk preferences, just that they use the same functional *form*. The studies differ in terms of how they use these bounds, as discussed briefly below. The use of bounds such as these is familiar from the laboratory experimental literature on risk aversion: see Holt and Laury [2002], Harrison, Johnson, McInnes and Rutström [2005] and Harrison, Lau, Rutström and Sullivan [2005] for discussion of how one can then use interval regression methods to analyse them. The limitation of this approach is that it is difficult to go beyond the CRRA or other one-parameter families, and in particular to examine other components of choice under uncertainty (such as more flexible utility functions, preference weighting or loss aversion).[38] Post, van den Assem, Baltussen and Thaler [2006] rely heavily on CRRA bounds in their analysis, and it has been employed in various forms by others as noted below.

The second is the examination of specific choices that provide "trip wire" tests of certain propositions of EUT, or provide qualitative indicators of preferences. For example, decisions made in the very last rounds often confront the contestant with the expected value of the unopened prizes, and allow one to identify risk lovers or risk averters directly. The limitations of this approach is that these choices are subject to sample selection bias, since risk attitudes and other preferences presumably played some role in whether the contestant reached these critical junctures. Moreover, they provide limited information at best, and do not allow one to define a metric for errors. If we posit some stochastic error specification for choices, as is now common, then one has no way of knowing if these specific choices are the result of such errors or a manifestation of latent preferences. Blavatskyy and Pogrebna [2006a][2006b] illustrate the sustained use of this type of empirical strategy, which is also used by other studies in some respects.

The third strategy it to propose a latent decision process and estimate the parameters of that process using maximum likelihood. This is the approach we favor, since it allows one to examine structural issues rather than rely on *ad hoc* proxies for underlying preferences. Our discussion of the literature therefore focuses on those that have used similar methods, and the technical similarities and differences to our approach.

Bombardini and Trebbi [2005]

The statistical analysis of Bombardini and Trebbi [2005; §3.1] (BT) is exemplary: it sets out a formal likelihood model built on a latent utility-maximizing choice process and solves for the entire path of choices.[39] In fact, their numerical methods involve solving the game by backward-induction

[38] Abdellaoui, Barrios and Wakker [2005] offer a one-parameter version of the EP function which exhibits non-constant RRA for empirically plausible parameter values. It does impose some restrictions on the variations in RRA compared to the two-parameter EP function, but is valuable as a parsimonious way to estimate non-CRRA specifications, and could be used for "bounds analyses" such as these.

[39] BT also consider the implications of allowing for the apparent fact that in the Italian version the banker knows what prizes are in the remaining, unopened boxes. The latter possibility raises some thorny estimation issues, since the contestant should view themselves in a strategic game under incomplete information with the banker. Thus an offer might affect their belief about the value of the prizes remaining in some complicated manner. Faced with relatively intractable strategic games in the analysis of the venerable

for each candidate preference parameter. So they only estimate the likelihood for the last 3 rounds of the game (p.14), to avoid the computational burden of solving for the larger game; our approach avoids this numerical complication, and allows us to estimate over the entire game.

BT assume a CRRA functional form for utility. They use this functional form to find the CRRA values at which the subject is indifferent between accepting the bank offer or not, and correctly use these to infer an interval response. They then specify the likelihood for an interval-censored model, recognizing that some subjects have closed CRRA intervals (those that accepted a bank offer) and others do not (e.g., those that do not accept any bank offer).[40] BT do not allow for any stochastic error in the decision-making process, although they do allow for sampling errors in the estimation stage. They consider three different, exogenous levels of wealth as arguments of the utility function. One is zero, one is a proxy for income, and one is a tenfold increase in the proxy for income and intended as a proxy for wealth. They also, like us, use a non-parametric bank offer function based on empirical data.

One limitation of their estimation method is that it does not correct for the possible correlation of errors by the same individual. In their model each contestant can have 1, 2 or 3 observations, depending on their choices. They note (p.16) three methods of allowing for taste heterogeneity, but none of these correct the estimates of standard errors for the *unobserved* individual heterogeneity common in panel models of this kind.[41]

BT consider differences in RRA when stakes differ (§4.4). They do so by comparing contestants that get into the last few rounds with low stakes remaining against contestants that get into the last few rounds with higher stakes remaining. They conclude that the former have much lower RRA than the latter, consistent with risk neutrality for low stakes and risk aversion for higher stakes. There are some obvious concerns that sub-sampling from later rounds might bias inferences, since subjects with higher aversion to risk would, *ceteris paribus*, have dropped out earlier. And variations in luck in earlier rounds would affect choices and sample selection into later rounds, as noted in our discussion of the effect of prior outcomes in Post, van den Assem, Baltussen and Thaler [2006b]. On the other hand, their conclusion is consistent with our conclusion using the entire game and a utility function that allows for varying RRA. One concern with these results of BT is that they would seem to invalidate the CRRA utility function they use throughout, raising specification issues.

A particularly interesting extension by BT is to estimate a non-EUT model (§5). They in fact estimate a RDEU model, or a version of CPT in the gain domain. Like us, they use a CRRA value function and a probability weighting function. Unfortunately they use an extremely restrictive functional form for the probability weighting specification, the power function $\omega(p) = p^{\varsigma}$. For values of the exponent $\varsigma<1$ ($\varsigma>1$), this implies overweighting (underweighting) *for all p*. Thus if the subjects

---

double-auction institution, theorists have often added simplifying assumptions to bypass such complications. For example, Friedman [1991] uses a Bayesian Game Against Nature assumption: that the agent behaves as if there is no strategic game with another sentient player, simply to avoid the cognitive burden of doing so. We focus on the version of their estimation that assumes that offers are not informative. BT conclude that it makes little difference to their estimates, which is consistent with that assumption. There is no evidence that the banker in the UK version knows the prizes.

[40] This is the interval-response method that would be appropriate for the CRRA-consistent bounds inferred by Post, van den Assem, Baltussen and Thaler [2006b].

[41] Their first method is to add a standard deviation around the estimate of the CRRA parameter *r*. Their second method is to allow for a small set of observable individual characteristics to affect the mean estimate of *r*. The set of characteristics they collect is larger than most in the DOND literature (see their Appendix C), but still small. Their third method is to allow for these observable characteristics to affect the estimated standard deviation of *r*, which is a heteroskedasticity correction.

exhibit probability weighting of the form we identify, consistent with the bulk of the RDEU and CPT literature, there would be a concave portion for small p and then a convex portion for larger p. In fact, if one estimates the conventional form (8) popularized by Tversky and Kahneman [1992] with our data, the function is concave up to p=0.37 and convex thereafter (the switch point *has* to be either 0.37 or 0.63). Hence, if probabilities for prizes in the last three rounds, the sample used by BT, are between 0.25 and 0.5, they would be forcing the power function to fit "data that wants to be concave and then convex." It is therefore not a surprise that they estimate ς=0.92 with a standard error of 0.06, and cannot reject the EUT-consistent null that ς=1. This is an artefact of assuming the restrictive power function, not an inference solely from the data.

Mulino, Scheelings, Brooks and Faff [2006]

The Australian version of *DOND* has several unique features. Two, in particular, "kick in" when a contestant is down to two or one unopened boxes after accepting an earlier bank offer. After the offer is accepted it is almost universal that the host counter-factually opens the remaining boxes, for various entertainment reasons: "did you make a good deal or not?" Depending on the history of the game, the host might play what is called a *Chance* round or a *SuperCase* round. In *Chance* the contestant can swap the accepted bank offer for a 50:50 chance at the remaining prizes. In *SuperCase*, which occurs when there is only one unopened case and all prizes are known, the contestant can swap the accepted bank offer for a lottery of 8 prizes in the low to middle range of the original prize set. These options are also used by the producers to "fill" an episode, so that there is no need to start a new contestant if the previous game finishes quickly.

Mulino, Scheelings, Brooks and Faff [2006] (MSBF) show that these options provide a form of insurance to contestants in earlier rounds, encouraging them to accept bank offers earlier than they might otherwise. They focus on the effect that these three different ways of framing the choice task affect elicited risk attitudes: the standard game, the choices in *Chance*, and the choices in *SuperCase*.

Their main statistical analysis consists of two parts. The first is a calculation of CRRA-consistent bounds for each subject in each round, in the same manner as Bombardini and Trebbi [2005] and Post, van den Assem, Baltussen and Thaler [2006]. These bounds are calculated allowing for a stochastic bank offer (as we do) as well as a deterministic bank offer, to show the effect of allowing for uncertainty about the bank offer. They also show the effect of allowing for the extra games, and that they indeed affect elicited risk attitudes.

The second part of their statistical analysis appears to be closer to our specification. They assume a latent CRRA utility function, a stochastic noise process in decision-making following the Luce specification used by Holt and Laury [2002]. However, they specify an equation (7) that sets the *decision* of the contestant (to deal or not) equal to the index of the ratio of the EU differences. This is conceptually the same as the term ∇EU in our equation (3), although the term is defined differently since we use a Fechner stochastic error specification instead of the Luce specification. This expression, in the form they define it, is already a cumulative density function, so they do not need the extra step (4) that we use; in this respect, their model follows Holt and Laury [2002]. But the right hand side of their equation (7) defines the *probability* of a given decision, not the binary variable identifying the *decision* itself. This is just a typographical error, and the functional form of the log-likelihood in their equation (8) is correct (Daniel Mulino; personal correspondence, 8/27/2006).

MSBF do not state (p.18) how they numerically evaluated the log-likelihood, and in particular the possible paths leading to the continuation values each subject faced. This is a non-trivial numerical issue: the methods we propose in §2 are intuitive but numerically intensive, and Bombardini and Trebbi [2005; p.14] find that they have to limit their analysis to just three final rounds for numerical reasons. The discussion of MSBF (§3.2.2, p.18) *suggests* that they used the intervals from the bounds calculations to provide starting values for a maximum-likelihood estimation (MLE), and then undertook a grid search within those bounds. This is also consistent with the explanation of the use of bootstrapping methods to identify standard errors for the MLE results (fn. 27). But the methods used to obtain MLE results are not clear, nor do they appear to be standard.

Post, van den Assem, Baltussen and Thaler [2006]

Interval Responses

The primary manner in which Post, van den Assem, Baltussen and Thaler (PABT) infer risk attitudes is by calculating CRRA-consistent bounds for each subject in each round. They do correct for the increasing bank offers in this calculation, taking into account the continuation value of turning down a bank offer in a given round. Consider the example in their Table III, described in PABT (p. 15). In round 1 this contestant has an upper bound CRRA of 6.94, since he would have to have a CRRA of this or higher to accept the deal given the options that he faced in future rounds (in particular, the improving bank offer). In round 2 his CRRA would have to have been 2.76 or higher to accept this deal.

One should then take these as interval responses, following the tabulation of CRRA intervals implied by observed lottery choices in Holt and Laury [2002; Table 3, p.1649] and the statistical analyses of Harrison, Johnson, McInnes and Rutström [2005] and Harrison, Lau, Rutström and Sullivan [2005]. Thus the first round for this subject implied an interval response of (-∞, 6.94], the second round implied an interval response of (-∞, 2.76], and so on. We do not *a priori* truncate the second round responses as (2.76, 6.94], since we want to remain open to finding shifts in risk attitudes over time, and collapsing the information from two rounds into one response would lose the original information on the two interval responses in different rounds. This is one of the hypotheses that PABT are interested in.

One might respond that the first round response is "uninformative" given the second round response. First, this misses the point about using interval responses directly since that is what the subject revealed: we could have used a later-round example to make that point. Second, while it is true that the first round response is not statistically informative with respect to the *mean* estimate of this subject, given the second round response, it is informative with respect to the *standard error* of that estimate. In other words, it is informative to know that the *same* subject has made *two* choices in which the latent value of a CRRA is 6.94 or lower instead of just *one* choice in which the latent value of a CRRA is 6.94 or lower.[42] This is particularly true for later rounds, but there is simply no rationale to discard the data from early rounds if one wants to estimate the standard error of the CRRA parameter in a full-information manner.

PABT do not want to calculate that standard error, at least from this stage of their statistical characterization. They want to construct bounds, based on the last two "active rounds" in the game for this subject. Thus, for the example in their Table III, the contestant accepted in round 5 where his CRRA switching value is 0.97, and his CRRA switching value on round 4 was 1.51. So this provides the bounds for their analysis: [0.97, 1.51]. In turn, they simply take an average of these bounds, on the assertion that "By construction, the upper and lower bounds are biased estimates. By averaging the positive and negative errors can be expected to cancel out, leading to a better estimate." (p.12). This is not quite correct. Their bounds are biased if one views them as the mean of the latent process generating the observed choices, and this seems to be what the intent is here. But

[42] One would have to correct for panel effects in this case, but such corrections are standard.

they are not biased if one views than as the "revealed preference" interval of values of some parameter of a latent process generating the observed choices. In this case the latent process, of course, is some individual behaving as if using a CRRA utility function and facing these opportunities. But if one has this latent process in mind, as PABT do, they should use the interval responses from each round and use standard interval-regression methods to infer the latent process, allowing for the panel nature of the data (since there are multiple observations per subject).[43]

Quite apart from the interval response being the correct unit of analysis given their assumptions about the latent decision process, there is a concern that the intervals they calculate are large. The example they give is quite large: [0.97, 1.51]. The problem is that even if one accepts their bounds, the imprecision here is 1.51 - 0.97 = 0.54 in CRRA units, and collapsing this to a mid-point of 1.24 with an imprecision of 0 clearly biases the analysis by introducing an error in the dependant variable. To be literal, [1.24, 1.24] is not the same as [0.97, 1.51]. Thus one does not know how many of their tests would survive when one recognizes that it is actually an interval between 0.97 and 1.5, rather than some precise point estimate (the average). In their equation (9) in PABT (p.20), used to test EUT, they simply put the inferred average on the left hand side as the dependent variable, ignoring the known noise around it from the interval.[44]

Prior Outcomes

PABT (p.21) devise a measure of the "fortune" of the player up to the last two "active rounds." This is just the ratio of the EV of the remaining prizes in that round to the EV of all prizes at the beginning of the game, and is a nice measure of the "joss," or good luck or bad luck, of the player in terms of opening prizes.[45] This measure, which we call joss, is then boiled down to a dummy Loss for losses (if the ratio is less than 1) or Gain for gains (if the ratio is greater than 1), and then used to generate measures of PriorLosses = (joss-1)×Loss and PriorGains = (joss-1)×Gain. We discuss the use of these measures below.

These measures play an important role in the main statistical inference of PABT, so it is worthwhile seeing what values they take. The pictures on the next page illustrate the density of PriorLosses (to the left of the horizontal axis and the vertical line at 0, in a solid line) and PriorGains (to the right of the horizontal axis, in a dashed line). These distributions use kernel densities for all subjects active in each round in the UK show. The distributions for country versions that last longer than the UK version are qualitatively the same, and illustrate the same points.[46]

---

[43] The estimation of these models has been a standard feature of popular statistical packages such as *Stata* or *LIMDEP* for many years. It is also relatively simply to code up directly as a specialized maximum likelihood routine.

[44] Footnote 4 (p.44) of PABT has some "corrections" to these steps when calculating the CRRA bounds, none of which are said to affect the main conclusions. First, the subject must have an EV of at least € 1000, or else they substitute the average bounds for subjects that stop in those rounds. Such subjects should probably just be dropped, since they are effectively out of the game. Second, subjects that do not accept a deal only have one end-point of the interval, so they assume that the other end-point is given by the round 9 end-point for the subject in question minus the average *spread* in end-points for all other subjects that stopped in the previous round. This issue can be more cleanly addressed if one moves to a statistical analysis directly defined over the interval-censored responses. In this case one just substitutes a clopen interval such as (-∞, 0] if the subject turned down a fair bank offer in the last round of choice (since a CRRA value of 0 indicates risk-neutrality).

[45] They take an average of this measure for the two rounds they consider, which probably does not make that much of a difference since EV is a martingale. If one were to keep the CRRA interval response from each round separate as a data point then it would be possible to use the value of this measure of luck from each round, which would be more precise.

[46] There is only one contestant with positive PriorGains in round 9 of the US show, so there is no display for that round.

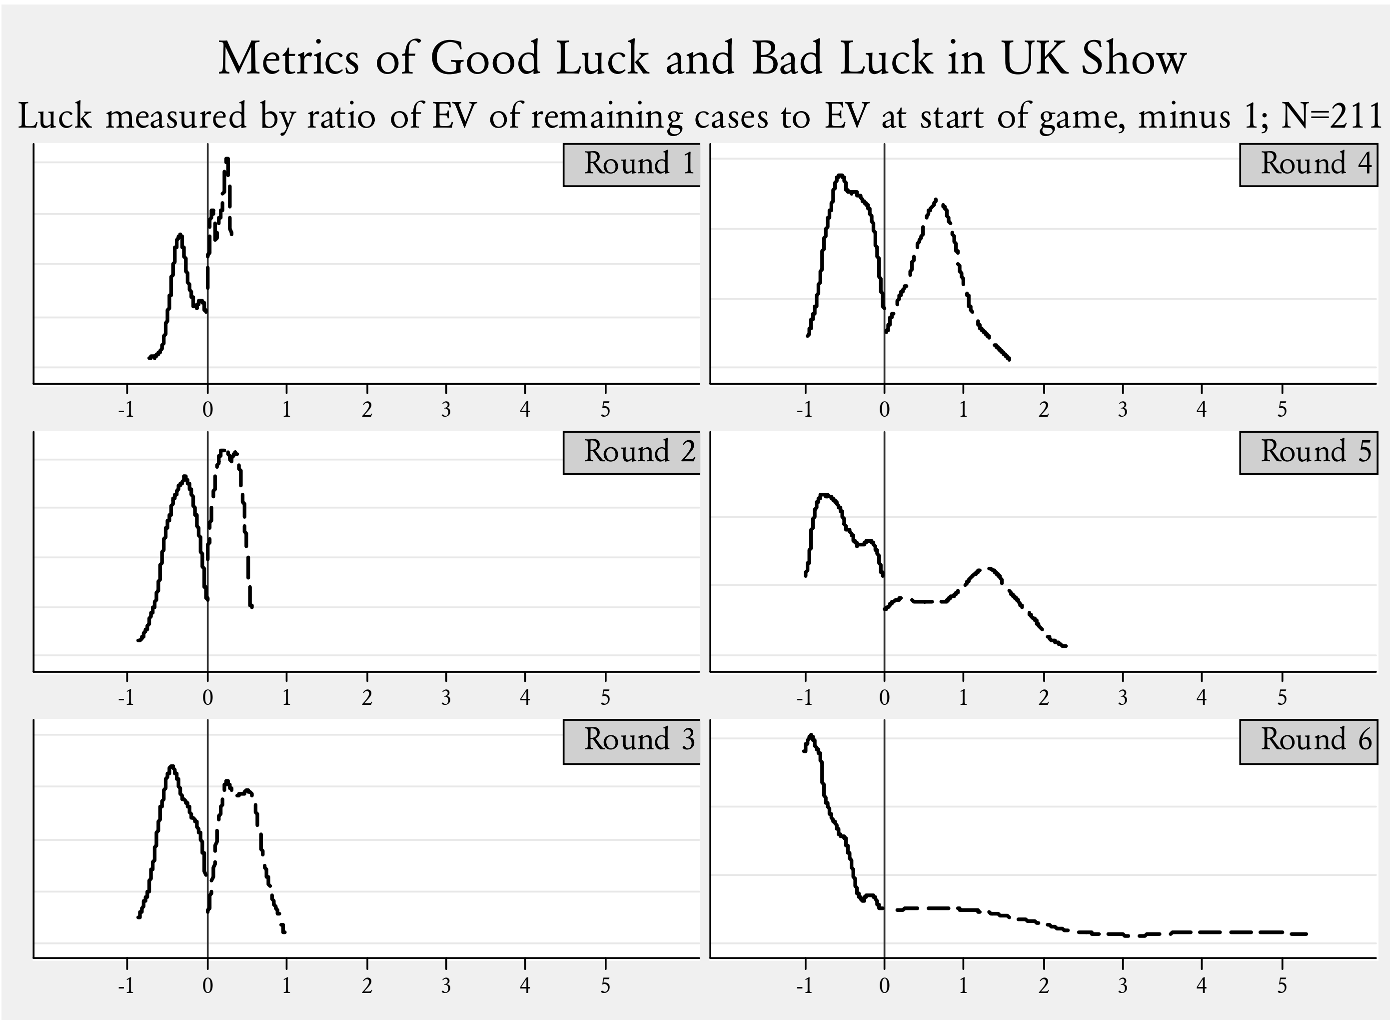
Metrics of Good Luck and Bad Luck in UK Show
Luck measured by ratio of EV of remaining cases to EV at start of game, minus 1; N=211
Round 1
Round 2
Round 3
Round 4
Round 5
Round 6
-1
0
1
2
3
4
5

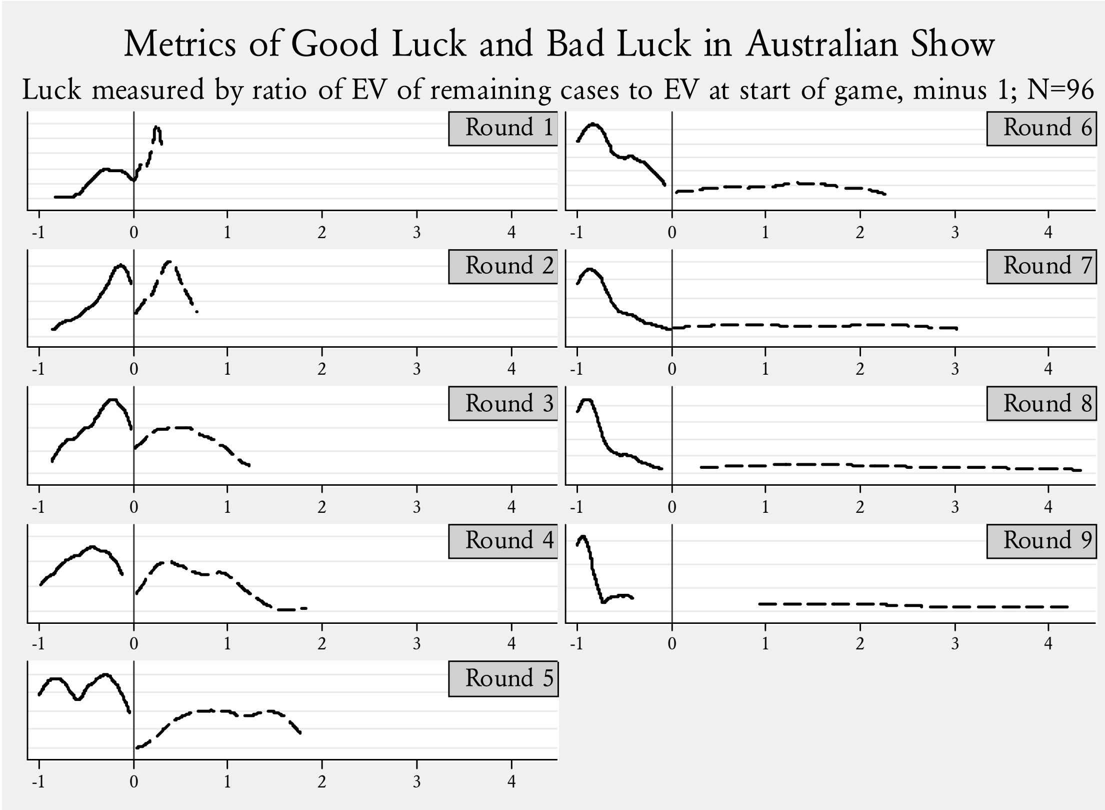
Metrics of Good Luck and Bad Luck in Australian Show
Luck measured by ratio of EV of remaining cases to EV at start of game, minus 1; N=96
Round 1
Round 2
Round 3
Round 4
Round 5
Round 6
Round 7
Round 8
Round 9
-1
0
1
2
3
4

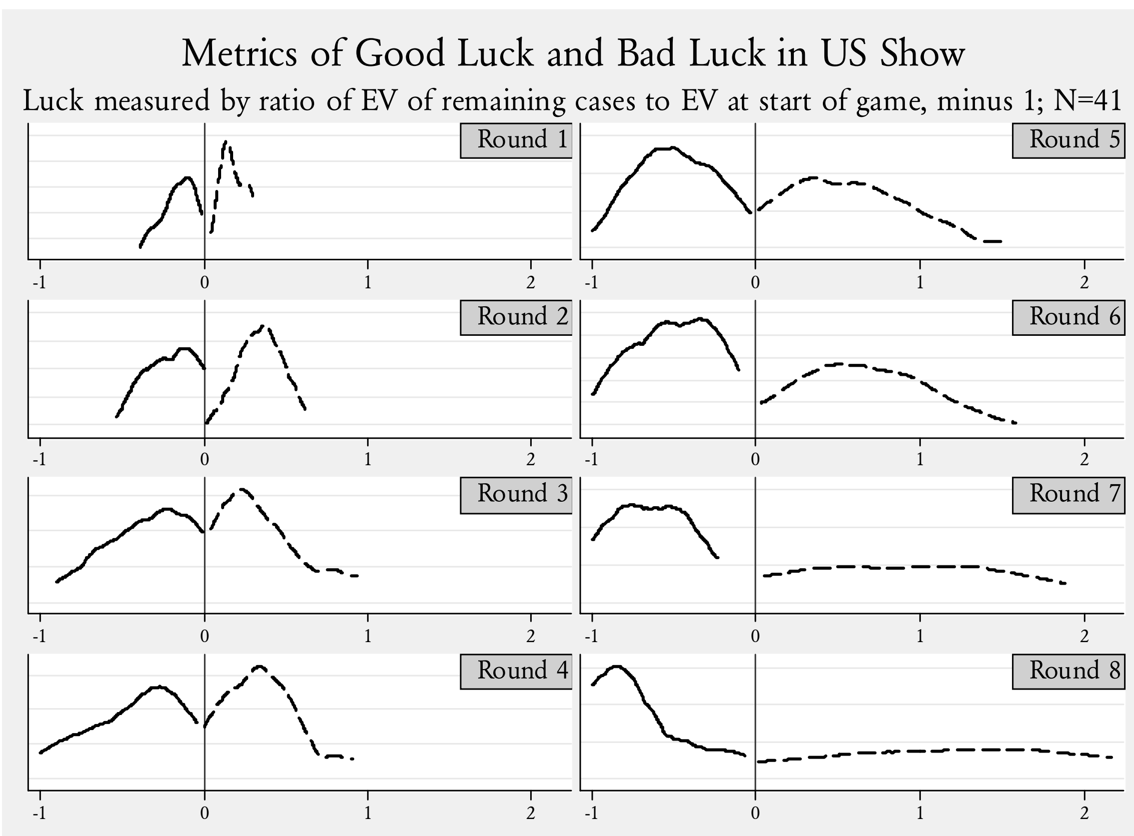
Metrics of Good Luck and Bad Luck in US Show
Luck measured by ratio of EV of remaining cases to EV at start of game, minus 1; N=41
Round 1
Round 2
Round 3
Round 4
Round 5
Round 6
Round 7
Round 8
-1
0
1
2

Focusing on the UK densities, consider the evolution of the measure as the game proceeds, recognizing that the sample changes as some contestants drop out. In round 1 the asymmetry in the two distributions is the direct result of an asymmetry in the *value* of the bottom half of the prizes and the top half of the prizes. In the UK version, as the earlier image of "Trevor" shows, there are just as many red prizes as blue prizes, but the *value* of the two sets is sharply asymmetric. Hence it is easier to suffer bad luck in EV terms than it is to enjoy good luck, and the distribution shows that. This asymmetry disappears as one proceeds into into rounds 2 and 3, and the measures are much more symmetric.

The fascinating evolution really starts in round 5 of the UK data. Consider the PriorGains measure first. It just collapses as one moves from round 4 into rounds 5 and 6, reflecting the attrrition of the sample of subjects that had good luck up to round 4. Their good luck translated into better offers, which induced them to accept the deal, *ceteris paribus* their risk attitudes. The other feature of the PriorGains measure is that it is not bounded at 1, whereas the PriorLosses measure is bounded at 1. So we start to see some subjects with very high values of this measure compared to the range of values of the PriorLosses measure. If all else were equal in terms of sample composition, and they are not, this would lead to larger standard errors on the PriorGains measure than the PriorLosses measure. This would make it *appear* as if bad luck is more important for behavior than good luck, but that is just an artefact of these being normalized differently. The implication is that the measures should be normalized, at the very least, if one wants to make inferences about their relative importance (and had some way of handling the sample attrition).

Turning to the PriorLosses measure, it changes from a symmetric, one-mode density in round 3 to become a severely positively-skewed density. This is again due to sample attrition, with the majority of "bad luck subjects" remaining in round 6 being those that had *severely* bad luck. In effect, their offers are so low that they reason that they may as well hope for the very best luck in the remaining stochastic realizations. Without knowing how such subjects change their reference point with the trail of bad luck they experience it is not possible to say if they are behaving in a loss averse or loss seeking manner.

The main concern is that the changes in these measures reflect sample attrition, which in turn reflects the core parameters one would like to estimate. This makes it very difficult to make simple claims about the effect of this measure on behavior, without a model of the choice process leading to the measure taking on certain values. That is, the measure is endogenous with respect to the parameters being measured, generating inconsistent estimates unless one accounts for the correlations implied by the endogeneity (in the standard manner of sample selection corrections).[47]

Stake Effects

To capture stake effects, or the possibility that RRA might not be constant, PABT introduce a dummy variable equal to the log of the EV at the outset of each game. This will pick up cross-country effects, since the stakes in the games are very different. Of course, it will also pick up differences in risk attitudes across countries, but they dismiss this as likely to be small compared to the pure stake effect: "Apart from the initial prizes, the editions used for our study are very similar and the contestants from the three European countries are comparable in terms of their cultural and economic background." (p.4)

Another concern is that there are in fact format differences across countries, as a further confound to the comparison of stakes. The Dutch *DOND* has 26 prizes and 9 rounds. However, the

[47] Our estimation approach avoids this problem by examining the entire history of each subject's choices in a game, and being explicit about the reference point. Of course, we assume that the reference point is exogenous, and that assumption deserves examination in future work.

German *DOND* consists of two versions. Only 6 of the German contestants in the PABT sample come from the version with 26 prizes and 9 rounds; the remaining 20 come from the version with 20 prizes and 8 rounds. These format differences may not be trivial in terms of the effect on the uncertainty of the lotteries implied by the game.

Finally, the sample sizes for the Belgian and German samples are very small: only 19 contestants and 26 contestants, respectively.

The reason that this issue is important is because it is the way in which PABT examine the possibility of varying RRA with stakes. They find (Table V, p.37) that a dummy for the size of the stakes in the cross-country editions is not statistically significant as a determinant of behavior, and use this to infer that there is no evidence to reject CRRA. Of course, this measure suffers from the confounds noted above, as well as the small samples (which bias hypothesis tests in favor of not finding significant effects if they are there for larger samples). Even without the confounds and small samples, it is a between-subjects test.

We examine this issue using an Expo-Power specification that looks at the wide range of prizes within a series in a given country. These differences are within-sample, since each subject faces a wide range of prizes in each round. This would appear to be a more reliable way to detect such differences.

Main Inferences

The main PABT results are presented in their Table V (p.37). We focus on the first column of results, assuming that the subjects do *not* integrate their prizes in *DOND* with their outside income or wealth, and that they rationally look forward in the game to improved bank offers.

The main effect, robust over the other variations, is that prior losses increases RRA. From a constant of 0.24 RRA is increased by 1.68, and this is statistically significant. The first concern is whether this survives recognition of the interval nature of the dependent variable.

The second concern is how this interacts with non-CRRA specifications, which our analysis finds to be important for the UK contestants. If someone has had bad luck, then by definition their lottery (whether myopic or forward-looking) is going to have more prizes that are of lower value than someone else. So if there is IRRA with prizes this IRRA will simply be picked up by the loss term. PABT would argue that this is controlled for by the "stake" variable, but that variable is confounded with other things noted above. In other words, PABT *might* be right that this is a pure prior loss effect, but there are confounds and the inference is not at all obvious from the analysis presented.[48]

Prospect Theory

Section VI is an innovative attempt to calibrate and estimate CPT for these data. PABT recognize (p.26) that CPT cannot be used to infer parameters for *each* subject as they did for CRRA and EUT, since there are too many parameters. But they can pool across subjects, and test CPT in

---

[48] PABT make an odd claim (p. 24) about the comparison of apples and oranges. Apparently they find that losers in large-stake countries are less risk averse than winners in small-stake countries. This is due to the difference in the size of the coefficients on the stake variable and the loss variable; but the stake variable is in logs, and measures EV across countries, and one does not know what size of average loss is captured by the loss variable. They then go on to compare winners in one country with losers in another country (Table VI), which potentially confounds several things and has very small samples. Again, the claims about these winners and losers might be correct, but these results are not convincing.

that manner. They also recognize some gaps in applying CPT here:

- CPT says nothing about multi-stage decisions. So they assume subjects are "myopic" in their value functions, although they define myopia as looking forward one period. This seems overly restrictive. The key point of CPT in this respect is not to integrate current income with past income, but that is a backwards-looking myopia rather than a forward-looking myopia.
- CPT does not specify how the reference point should change in dynamic games. But Thaler & Johnson [1990] do have a good discussion of some alternatives, and *DOND* suggests some obvious ones: zero, the last bank offer, or the highest bank offer so far. One additional possibility is to use the average earnings in the broadcast show up to that point for the opening round, until a bank offer exceeds that. Given the popularity of the show, it can be expected after several broadcasts that contestants have a good sense of the rough level of expected earnings in these games, and what is a good outcome and a bad outcome historically.

The estimation procedure used by PABT is then to find parameters that allow the "hit rate" for CPT to exceed the hit rate for EUT. They find some parameters (p.28) that take the hit rate to 65% or 69%, and declare victory over EUT, which had hit rates of 60% or 61%. Of course, this comparison of hit rates contains no basis for statistical comparison, since we have no measure of the imprecision of the estimates underlying the hit rates. This is precisely what likelihoods do in a complete statistical analysis.

Moreover, one would ideally want to characterize two latent processes here, EUT and CPT, and allow the data to identify the relative weights on each. Of course, there are several EUT variants (e.g., CRRA, expo-power) and several CPT variants (e.g., reference points, different probability weighting functions). But one can presumably boil these down to some interesting alternatives to compare, since the data requirements for more than two processes can be daunting.

De Roos and Sarafidis [2006]

The Australian version of DOND provides the data for De Roos and Sarafidis [2006] (DS), who focus on the estimation of risk attitudes assuming EUT as well as a RDEU specification. They employ two statistical methods: an exploratory CARA-consistent bounds analysis, and then a full-blown maximum likelihood model initially using CRRA and CARA. There are two important methodological features of the latter model: it deals with panel effects rigorously, and is then extended to a RDEU specification. The bounds analysis is viewed as a relatively non-parametric way of motivating the need to worry about individual (observed and unobserved) heterogeneity.

The ML model proposed by DS allows for noise in the bank offer, but their empirical implementation does not include it. However, they do allow for a stochastic noise component in the latent decision-making process. This specification is nicely motivated (p.12) by noting that a number of subjects appear to have behaved inconsistently with EUT, at least if one assumes CARA. In each case the inconsistency involved a reversal of risk attitudes as rounds progressed: someone accepting an offer that had been revealed as worse than a rejected offer in a previous round, for example. Although one would want to account for stochastic errors consistently with estimates of the sampling error of the risk preferences before declaring such choices EUT-inconsistent, these motivating examples are correct for this pedagogic purpose.

The methodological highlight of the base EUT specification of DS is the definition of three ways of allowing for individual heterogeneity (p.17). One is to ignore it. The other is to assume a random effects term ν for each individual and add it to the latent index defining the probability of choosing deal. This is changing our specification

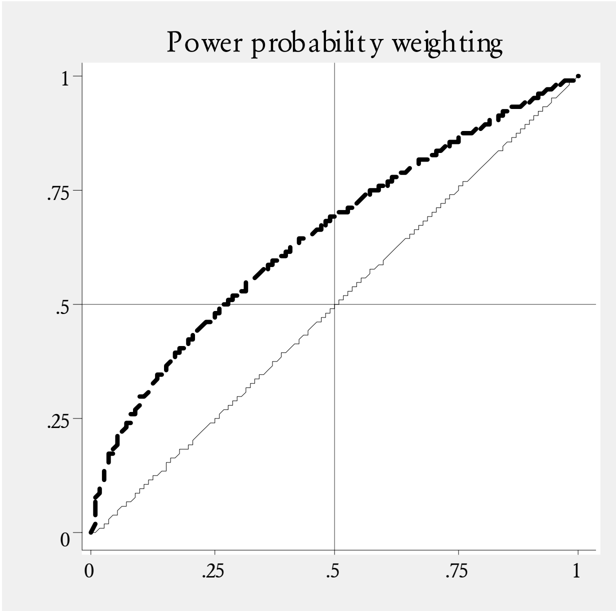

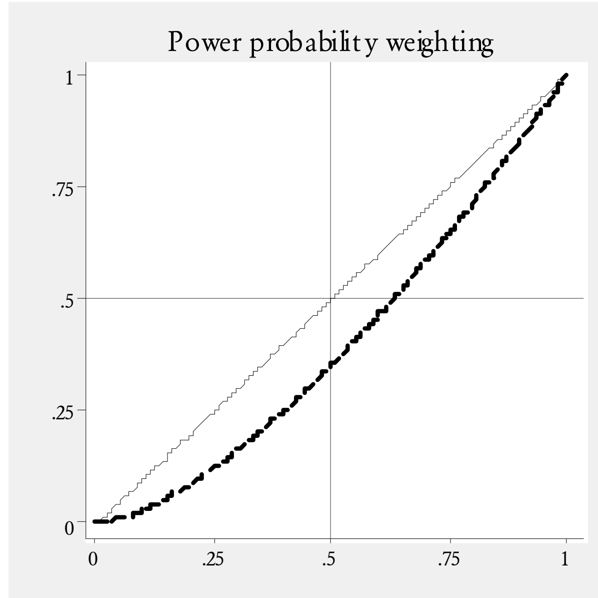

$$G(\nabla EU) = \Phi(\nabla EU). \quad (4)$$

to

$$G(\nabla EU) = \Phi(\nabla EU) + \nu. \quad (4')$$

The final method is to view the CRRA coefficient as a random coefficient reflecting a subject specific random effect υ, so that our specification

$$\hat{r} = \hat{r}_0 + (\hat{r}_{FEMALE} \times FEMALE) + (\hat{r}_{YOUNG} \times YOUNG) + (\hat{r}_{OLD} \times OLD) \quad (6)$$

would become

$$\hat{r} = \hat{r}_0 + (\hat{r}_{FEMALE} \times FEMALE) + (\hat{r}_{YOUNG} \times YOUNG) + (\hat{r}_{OLD} \times OLD) + \upsilon. \quad (6')$$

In each case the random effect term ν and υ varies *across* individuals but is the same for each choice by the *same* individual. These are valid statistical ways to allow for individual heterogeneity, but tend to be computationally intensive.

The second methodological highlight of DS is their extension to consider a RDEU specification. They do so by assuming that decisions are made in the gain frame, as we do, and as do Bombardini and Trebbi [2005; §5] discussed above. Unfortunately they also employ the restrictive power probability weighting function $\omega(p) = p^{\varsigma}$, which restricts one to estimate either (weakly) concave or (weakly) convex decision weights for the entire range of probabilities. They estimate $\varsigma \approx$ ½ using CRRA and the dynamic (forward-looking) model with corrections for individual heterogeneity. This implies severe over-weighting of probabilities, as shown to the left above.

One "extension" of their RDEU specification is to estimate the early RDEU model of Yaari [1987]. This amounts to assuming that there is no diminishing marginal utility of income, and that contestants are risk neutral. Hence all of the explanatory power rests on the probability weighting function. In this case they estimate $\varsigma \approx 1½$, implying the under-weighting for all p shown to the right above.

Other Studies

Deck, Lee and Reyes [2006] estimate CRRA-consistent and CARA-consistent bounds from observed choices in the Mexican version of *DOND*. They use the last few rounds, since the computational burden of evaluating every possible continuation path became too large for earlier rounds (p.6). Our numerical approach, on the other hand, samples from these paths in a Monte Carlo manner, so does not become numerically crippled when evaluating the whole game. Of course, this means that there is some approximation error from the Monte Carlo sampling.

Blavatskyy and Pogrebna [2006a][2006b] do not undertake estimation of a model of the preferences of any latent decision-making process. Instead, they seek to identify "trip wire" tests of certain theoretical predictions. The problem with this approach is that is provides no metric for stochastic error: a "small deviation" from the prediction of a deterministic model counts as much as a "large deviation." The lack of any consideration of stochastic error is odd, given the important contribution of Blavatsky [2005] on exactly this general issue. Blavatskyy and Pogrebna [2006c] pick up this point, by focusing solely on the role of stochastic errors in formal models of *DOND*. Unfortunately their models ignore the dynamic aspects of the game, assuming that contestants simply compare the current bank offer with the single virtual lottery that involves the subject saying "No Deal" to *every* future bank offer. This turns out to be reasonable approximation for the UK, as we show in Figure 2, but in general cannot be relied upon.

Botti, Conte, DiCagno and D'Ippoliti [2006] also focus on a static representation of the game. They follow De Roos and Sarafidis [2006] and consider the role of unobserved individual heterogeneity in behavior, using techniques from maximum simulated likelihood.

Additional References

Blavatskyy, Pavlo, "A Stochastic Expected Utility Theory," *IEW Working Paper No. 231*, Institute for Empirical Research in Economics, University of Zurich, February 2005.

Friedman, Daniel, "A Simple Testable Model of Double Auction Markets," *Journal of Economic Behavior & Organization*, 15, 1991, 47-70.

Harrison, Glenn W.; Johnson, Eric; McInnes, Melayne M., and Rutström, E. Elisabet, "Risk Aversion and Incentive Effects: Comment," *American Economic Review*, 95(3), June 2005, 897-901.

Harrison, Glenn W.; Lau, Morten Igel; Rutström, E. Elisabet, and Sullivan, Melonie B., "Eliciting Risk and Time Preferences Using Field Experiments: Some Methodological Issues," in J. Carpenter, G.W. Harrison and J.A. List (eds.), *Field Experiments in Economics* (Greenwich, CT: JAI Press, Research in Experimental Economics, Volume 10, 2005).