

Dual Criteria Decisions

Andersen, Steffen; Harrison, Glenn W.; Lau, Morten I.; Rutström, E. Elisabet

Document Version
Final published version

Publication date:
2009

License
CC BY-NC-ND

Citation for published version (APA):
Andersen, S., Harrison, G. W., Lau, M. I., & Rutström, E. E. (2009). *Dual Criteria Decisions*. Department of Economics. Copenhagen Business School. Working Paper / Department of Economics. Copenhagen Business School No. 2-2009

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 30. Mar. 2023





**Copenhagen
Business School**
HANDELSHØJSKOLEN

Department of Economics

Copenhagen Business School

Working paper 2-2009

DUAL CRITERIA DECISIONS

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

Dual Criteria Decisions

by

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

June 2007

Working Paper 06-11, Department of Economics,
College of Business Administration, University of Central Florida, 2006

Abstract. The most popular models of decision making use a single criteria to evaluate projects or lotteries. However, decision makers may actually consider multiple criteria when evaluating projects. We consider a dual criteria model from psychology. This model integrates the familiar tradeoffs between risk and utility that economists traditionally assume, allowance for rank-dependent decision weights, and consideration of income thresholds. We examine the issues involved in full maximum likelihood estimation of the model using observed choice data. We propose a general method for integrating the multiple criteria, using the logic of mixture models, which we believe is attractive from a decision-theoretic and statistical perspective. The model is applied to observed choices from a major natural experiment involving intrinsically dynamic choices over highly skewed outcomes. The evidence points to the clear role that income thresholds play in such decision making, but does not rule out a role for tradeoffs between risk and utility or probability weighting.

[†] Department of Economics and 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, UK (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.

When decisions are being made about risky investments, do decision makers boil all of the facets of the prospect down to one criterion, which is then used to rank alternatives and guide choice, or do they use multiple criteria? The prevailing approach of economists to this problem is to assume a single criterion, whether it reflects standard expected utility theory (EUT), rank-dependent utility (RDU) theory, or prospect theory (PT). In each case the risky prospect is reduced to some scalar, representing the preferences, framing and budget constraints of the decision-maker, and then that scalar is used to rank alternatives.¹

Many other disciplines assume the use of decision-making models with multiple criteria.² In some cases these models can be reduced to a single criterion framework, and represent a recognition that there may be many attributes or arguments of that criteria.³ And in some cases these criteria do not lead to crisp scalars derivable by formulae.⁴ But often one encounters decision rules which provide different metrics for evaluating what to do, or else one encounters frustration that it is not possible to encapsulate all aspects of a decision into one of the popular single-criteria models.

We consider this “dual criteria” approach by means of an extraordinarily rich case study: the television game show *Deal Or No Deal*. Behavior in this show provides a wonderful opportunity to

¹ In economics the only exceptions are lexicographic models, although one might view the criteria at each stage as being contemplated simultaneously. For example, Rubinstein [1988] and Leland [1994] consider the use of similarity relations in conjunction with “some other criteria” if the similarity relation does not recommend a choice. In fact, Rubinstein [1988] and Leland [1994] reverse the sequential order in which the two criteria are applied, indicating some sense of uncertainty about the strict sequencing of the application of criteria. Similarly, the original prospect theory of Kahneman and Tversky [1979] considered an “editing stage” to be followed by an “evaluation stage,” although the former appears to have been edited out of later variants of prospect theory.

² Quite apart from the model from psychology evaluated here, there is a large literature in psychology referenced by Starmer [2000] and Brandstätter, Gigerenzer and Hertwig [2006]. Again, many of these models of the decision process present multiple criteria that might be used in a strict sequence, but which are sometimes viewed as being used simultaneously. In decision sciences the weighted sum model of Fishburn [1967] remains popular, although it could be viewed as a multi-attribute utility model. The analytic hierarchy process model of Saaty [1980] remains very popular in corporate settings, and has gone through numerous revisions and extensions. Popular textbooks on multi-criteria decision making in business schools include Kirkwood [1997] and Liberatore and Nydick [2002]; the emphasis at that level is on alternative software packages that are commercially available. There also exists a *Journal of Multi-Criteria Decision Analysis* (<http://www3.interscience.wiley.com/cgi-bin/jhome/5725>).

³ For example, multi-attribute expected utility, reviewed in Keeney and Raiffa [1976] or von Winterfeldt and Edwards [1986; ch.7]. Or one can seek appropriate single-criteria utility representations of informal dual-criteria decision rules, such as the well-known tradeoff between “risk” and “return” (e.g., Bell [1995]).

⁴ For example, old debates in psychology about when one should use “heads instead of formulas,” reviewed by Kleinmütz [1990]. Also see Hogarth [2001] for a related perspective.

examine dynamic choice under uncertainty in a controlled manner with substantial stakes. The show has many of the features of 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.⁵

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.⁶ We explain the format of the show in section 1, and discuss this temporal connection.

We examine two modeling approaches to these data. One is the single-criterion RDU model, which can be viewed as a generalization of EUT to allow for non-linear decision weights. The other is a dual-criteria model from psychology which could have been built with this task domain in mind: the SP/A theory of Lopes [1995]. The SP/A model departs from EUT, RDU and PT in one major respect: it is a dual criteria model. Each of the single criteria models, even if they have a number of components to their evaluation stage, boil down to a scalar index for each lottery. The SP/A model instead explicitly posits two distinct but simultaneous ways in which the same subject might evaluate a given lottery. One is the SP part, for a process that weights the “security” and “potential” of the

⁵ Game shows are increasingly recognized as a valuable source of replicable data on decision-making with large stakes. Andersen, Harrison, Lau and Rutström [2007] review the applications to choice under uncertainty, including the many recent applications of data from *DOND*. All of these studies consider single-criteria models.

⁶ 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.”

lottery in ways that are similar to RDU. The other is the A part, which focuses on the “aspirations” of the decision-maker. In many settings these two parts appear to be in conflict, which means that one must be precise as to how that conflict is resolved. We discuss each part, and then how the two parts may be jointly estimated in section 2.

Apart from presenting a systematic maximum-likelihood approach to the estimation of the SP/A model, we propose a natural decision-theoretic and statistical framework to resolve the potential conflict between the two criteria. This is the notion of a mixture of latent decision-making processes. Rather than view the observed data as generated by a single decision-making process, such as EUT, RDU *or* PT, one could easily imagine the data from a sample being generated by some mixture of these processes. Harrison and Rutström [2005], for example, allowed (laboratory lottery) choices to be made by EUT *and* PT, with a statistical mixture model being used to estimate the fraction of choices better characterized by EUT and the fraction better characterized by PT. In our case we simply extend this mixture notion to the two criteria of one model, rather than the two criteria of two models. We discuss this approach, and its interpretation, in section 3. We argue that mixture models provide a natural formalization, in theory and applied work, of multiple-criteria models.

We present empirical results in section 4, estimating an RDU model and then an SP/A model with data drawn from the UK version of *Deal Or No Deal*. We employ data covering 2,317 choices by 461 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 £17,737 in our sample. The distribution of earnings is heavily skewed, with relatively few subjects receiving the highest prizes, and median earnings are £13,000.

We find evidence that there is indeed some probability weighting being undertaken by contestants. We also find evidence that “aspiration levels” and “security levels” play a role in decision-making in the SP/A model, which was motivated by psychological findings in task domains that have highly skewed prize distributions. To some extent one can view these aspiration and security levels as similar to reference points and loss aversion, concepts from PT, although the

psychological motivation and formal modeling is quite distinct. Thus we conclude that more attention should be paid to the manner in which psychologically-motivated notions of choice in risky behavior are modeled.

In summary, in section 1 we document the game show format and field data we use. In section 2 we describe the general statistical models developed for these data, assuming an SP/A model of the latent decision-making process. In section 3 we review the use and interpretation of mixture specifications in dual criteria models. Section 4 presents empirical results from estimating the SP/A model using the large-stakes game show data, and section 5 examines implications for model comparisons. Finally, section 6 offers 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 boxes 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 rounds 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 through 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.⁷

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 earlier rounds. In our data, only 39% of contestants reach that point.⁸ 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

⁷ 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 have this feature generally, making the analysis much cleaner.

⁸ 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. In the UK version the generosity of later-round bank offers slowly improved over the seasons of the show, and we allow for this by using a lagged estimate of the empirical distribution of offers.

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 three more boxes are 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.

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.⁹ 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 expected utility of the compound lottery. This corresponds to procedures developed in the finance literature to price path-dependent 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),¹⁰ that the dollar

⁹ 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.

¹⁰ 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

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/>. Our data consists of behavior on 461 contestants.

2. Modeling Contestant Behavior

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 of 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 (convex) probability weighting functions would imply risk seeking (risk aversion). 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.¹¹ In economics Quiggin [1982] 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).

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.

¹¹ 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 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 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. We return below to consider a later formalization of the ideas of Lopes [1984].

In the RDU model utility can be defined over money m using a Constant Relative Risk Aversion (CRRA) function

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

where $\rho \neq 1$ is the RRA coefficient, and $u(m) = \ln(m)$ for $\rho = 1$. With this parameterization, and assuming EUT, $\rho = 0$ denotes risk neutral behavior, $\rho > 0$ denotes risk aversion, and $\rho < 0$ denotes risk loving. In fact, under RDU there is more to the characterization of risk attitudes than the concavity of the utility function, so we refer instead to ρ as simply controlling the curvature of the utility function, rather than defining risk attitudes.

Let p_k denote the probability induced by the task for outcome k . To calculate decision weights under RDU one replaces the expected utility of lottery i ,

$$EU_i = \sum_{k=1, 20} [p_k \times u_k], \quad (2)$$

with the rank-dependent utility of the lottery,

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

where

$$w_i = \omega(p_i + \dots + p_n) - \omega(p_{i+1} + \dots + p_n) \quad (3a)$$

for $i=1, \dots, n-1$, and

$$w_i = \omega(p_i) \quad (3b)$$

for $i=n$, where the subscript indicates outcomes ranked from worst to best, and $\omega(p)$ 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. It is assumed to have well-behaved endpoints such that $\omega(0)=0$ and $\omega(1)=1$ and to imply weights

$$\omega(p) = p^\gamma / [p^\gamma + (1-p)^\gamma]^{1/\gamma} \quad (4)$$

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,” 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 *overweighting* 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.¹² Under RDU $\gamma > 1$ implies *underweighting* of extreme outcomes.

The rank-dependent utility of the bank offer, RDU_{BO} , can be evaluated directly from (1) since it involves no uncertainty. It can then be compared to the rank-dependent utility of the lottery induced by saying “No Deal,” RDU_{ND} , and allowance made for a Fechner “noise” term μ to accommodate the possibility that the decision-maker makes some errors when comparing these scalars. Thus we have the latent index

$$\nabla RDU = (RDU_{BO} - RDU_L) / \mu \quad (5)$$

to show the strength of the latent preference for the bank offer, and saying “Deal” in this round. This latent index can be transformed into a probability of saying “Deal” or “No Deal” using a

¹² There are some well-known limitations of the probability weighting function (4). 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 . There exist two-parameter probability weighting functions that exhibits more flexibility than (4), but for our purposes the standard probability weighting function is adequate.

cumulative density function $G(\nabla RDU)$, and the conditional log-likelihood becomes

$$\ln L^{RDU}(\rho, \gamma, \mu; y, X) = \sum_i l_i^{RDU} = \sum_i [(\ln G(\nabla RDU) | y_i=1) + (\ln (1-G(\nabla RDU)) | y_i=0)] \quad (6)$$

where y is the observed choice of “Deal” ($y=0$) or “No Deal” ($y=1$). This likelihood requires the estimation of ρ , γ and μ . We view it as conditional in the sense that it assumes the RDU model of the latent decision process, as well as the parametric forms (1), (4) and (5).

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

$$RDEV_i = \sum_{k=1,20} [w_k \times m_k] \quad (2'')$$

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

B. Rank and Sign-Dependent Preferences: SP/A Theory

Kahneman and Tversky [1979] introduced the notion of sign-dependent preferences, stressing the role of the reference point when evaluating lotteries. The notion of rank-dependent decision weights was incorporated into their sign-dependent PT by Starmer and Sugden [1989], Luce and Fishburn [1991] and Tversky and Kahneman [1992]. Unfortunately, economists tend to view psychological models as monolithic, represented by the variants of PT. In fact there are many alternative models in the literature, although often they have not been developed in a way that would facilitate application and estimation.¹³ One that seems unusually well suited to the *DOND* environment is also rank and sign-dependent: the SP/A theory of Lopes [1995].

The SP/A model departs from EUT, RDU and PT in one major respect: it is a dual criteria model. Each of the single criteria models, even if they have a number of components to their evaluation stage, boil down to a scalar index for each lottery such as (2), (2') and (2''). The SP/A model instead explicitly posits two distinct but simultaneous ways in which the same subject might evaluate a given lottery. One is the SP part, for a process that weights the “security” and “potential”

¹³ Starmer [2000] provides a well-balanced review from an economist’s perspective.

of the lottery in ways that are similar to RDEV. The other is the A part, which focuses on the “aspirations” of the decision-maker. In many settings these two parts appear to be in conflict, which means that one must be precise as to how that conflict is resolved. We discuss each part, and then how the two parts may be jointly estimated.

Although motivated differently, the SP criteria is formally identical to the RDEV criteria reviewed earlier. The decision weights in SP/A theory derive from a process by which the decision-maker balances the security and potential of a lottery. On average, the evidence collected from experiments, such as those described in Lopes [1984], seems to suggest that an inverted-S shape familiar from PT

... represents the weighting pattern of the average decision maker. The function is security-minded for low outcomes (i.e., proportionally more attention is devoted to worse outcomes than to moderate outcomes) but there is some overweighting (extra attention) given to the very best outcomes. A person displaying the cautiously hopeful pattern would be basically security-minded but would consider potential when security differences were small. (Lopes [1995; p.186])

The upshot is that the probability weighting function

$$\omega(p) = p^\gamma / [p^\gamma + (1-p)^\gamma]^{1/\gamma} \quad (4)$$

from RDU would be employed by the *average* subject, with the expectation that $\gamma < 1$.¹⁴ However, there is no presumption that any individual subject follow this pattern. Most presentations of the SP/A model assume that subjects use a linear utility function, but this is a convenience more than anything else. Lopes and Oden [1999; p.290] argue that

Most theorists assume that [utility] is linear without asking whether the monetary range under consideration is wide enough for nonlinearity to be manifest in the data. We believe that [utility] probably does have mild concavity that might be manifest in some cases (as, for example, when someone is considering the huge payouts in state lotteries). But for narrower ranges, we prefer to ignore concavity and let the decumulative weighting function carry the theoretical load.

So the SP part of the SP/A model collapses to be the same as RDEU, although the interpretation of the probability weighting function and decision weights is quite different. Of course, the stakes in

¹⁴ Lopes and Oden [1999; equation (10), p.290] propose an alternative function which would provide a close approximation to (4). Their function is a weighted average of a convex and concave function, which allows them to interpret the average inverted S-pattern in terms of a weighted mixture of security-minded subjects and potential-minded subjects.

DOND are huge, so it is appropriate to allow for non-linear utility. Thus we obtain the likelihood of the observed choices conditional on the SP criteria being used to explain them; the same latent index (5) is constructed, and the likelihood is then (6) as with RDEU. The typical element of that log-likelihood for observation i can be denoted l_i^{SP} .

The aspiration part of the SP/A model collapses the indicator of the value of each lottery down to an expression showing the extent to which it satisfies the aspiration level of the contestant. This criterion is sign-dependent in the sense that it defines a threshold for each lottery: if the lottery exceeds that threshold, the subject is more likely to choose it. If there are up to K prizes, then this indicator is given by

$$A_i = \sum_{k=1, K} [\eta_k \times p_k] \quad (7)$$

where η_k is a number that reflects the degree to which prize m_k satisfies the aspiration level.

Although Oden and Lopes [1997] advance an alternative interpretation using fuzzy set theory, so that η_k measures the degree of membership in the set of prizes that are aspired to, we can view this as simply a probability. It could be viewed as a crisp, binary threshold for the *individual* subject, which is consistent with it being modelled as a smooth, probabilistic threshold for a *sample* of subjects, as here.

This concept of aspiration levels is close to the notion of a threshold income level debated by Camerer, Babcock, Loewenstein and Thaler [1997] and Ferber [2005]. The concept is also reminiscent of the “safety first” principle proposed by Roy [1952][1956] and the “confidence limit criterion” of Baumol [1963], although in each case these are presented as extensions of an expected utility criteria rather than as alternatives. It is also related to the vast literature on “chance-constrained programming,” applied to portfolio issues by Byrne, Charnes, Cooper and Kortanek [1967][1968].

The implication of (7) is that one has to estimate some function mapping prizes into probabilities, to reflect the aspirations of the decision-maker. We use an extremely flexible function for this, the cumulative non-central Beta distribution defined by Johnson, Kotz and Balakrishnan [1995]. This function has three parameters, χ , ξ and ψ . We employ a flexible form simply because

we have no *a priori* restrictions on the shape of this function, other than those of a cumulative density function, and in the absence of theoretical guidance prefer to let the data determine these values.¹⁵ We want to allow it to be a step function, in case the average decision-maker has some crisp focal point such as £25,000, but the function should then determine the value of the focal point (hence the need for a non-central distribution, given by the parameter ψ). But we also want to allow it to have an inverted S-shape in the same sense that a logistic curve might, or to be convex or concave over the entire domain (hence the two parameters χ and ξ).

Once we have values for η_k it is a simple matter to evaluate A_i using (7). We then construct the likelihood of the data assuming that this criteria was used to explain the observed choices. The likelihood, conditional on the A criterion being the one used by the decision maker, and our functional form for η_k , depends on the estimates χ , ξ , ψ and μ given the above specification and the observed choices. The conditional log-likelihood is

$$\ln L^A(\chi, \xi, \psi, \mu; y) = \sum_i l_i^A = \sum_i [(\ln G(\nabla A) \mid y_i=1) + (\ln (1-G(\nabla A)) \mid y_i=0)] \quad (8)$$

in the usual manner.

3. Mixtures of Decision Criteria

There is a deliberate ambiguity in the manner in which the SP and A criteria are to be combined to predict a specific choice. One reason is a desire to be able to explain evidence of intransitivities, which figures prominently in the psychological literature on choice (e.g., Tversky [1969]). Another reason is the desire to allow context to drive the manner in which the two criteria are combined, to reconcile the model of the choice process with evidence from verbal protocols of decision makers in different contexts. Lopes [1995; p.214] notes the SP/A model can be viewed as a function F of the two criteria, SP and A, and that it

... combines two inputs that are logically and mathematically distinct, much as Allais [1979] proposed long ago. Because SP and A provide conceptually independent assessments of a gamble's attractiveness, one possibility is that F is a weighted

¹⁵ It also helps that this function can be evaluated as an intrinsic function in advanced statistical packages such as *Stata*.

average in which the relative weights assigned to SP and A reflect their relative importance in the current decision environment. Another possibility is that F is multiplicative. In either version, however, F would yield a unitary value for each gamble, in which case SP/A would be unable to predict the sorts of intransitivities demonstrated by Tversky [1969] and others.

These proposals involve creating a unitary index of the relative attractiveness of one lottery over another, of the form

$$\nabla_{SP/A} = [\theta \times \nabla_{SP}] + [(1-\theta) \times \nabla_A] \quad (9)$$

for example, where θ is some weighting constant that might be assumed or estimated.¹⁶ This scalar measure might then be converted into a cumulative probability $G(\nabla_{SP/A}) = \Phi(\nabla_{SP/A})$ and a likelihood function defined.

A more natural formulation is provided by thinking of the SP/A model as a mixture of two distinct latent, data-generating processes. If we let π^{SP} denote the probability that the SP process is correct, and $\pi^A = (1-\pi^{SP})$ denote the probability that the A process is correct, the grand likelihood of the SP/A process as a whole can be written as the probability weighted average of the conditional *likelihoods*. Thus the likelihood for the overall SP/A model is defined as

$$\ln L(\rho, \gamma, \chi, \xi, \psi, \mu, \pi^{SP}; y, X) = \sum_i \ln [(\pi^{SP} \times I_i^{SP}) + (\pi^A \times I_i^A)]. \quad (10)$$

This log-likelihood can be maximized to find estimates of the parameters of each latent process, as well as the mixing probability π^{SP} . The literal interpretation of the mixing probabilities is at the level of the observation, which in this instance is the choice between saying “Deal” or “No Deal” to a bank offer. In the case of the SP/A model this is a natural interpretation, reflecting two latent psychological processes for a given contestant and decision.¹⁷

¹⁶ Lopes and Oden [1999; equation 16, p.302] offer a multiplicative form which has the same implication of creating one unitary index of the relative attractiveness of one lottery over another.

¹⁷ Byrne et al. [1967; p.19] elegantly view the multiple-criteria problem as characterizing the *objective of the latent decision-maker* as a probability distribution rather than reducing it to a scalar: “Some of the approaches we shall examine are also concerned with choices that maximize a single figure of merit. Others are concerned with developing the relevant combinations of probability distributions so that these may themselves be used as a basis for managerial choice. [...] To avoid misunderstanding it should be said, at this point, that this paper is not concerned with issues such as whether a ‘present value’ provides a better figure of merit than an ‘internal rate of return’ via a ‘bogey adjustment’ or a ‘payback period’ computation. Indeed it will be one purpose of this paper to suggest that some of these issues might be resolved – or at least placed in a different perspective – if some of the new methodologies can make it possible to avoid insisting on the use of one of these figures to the exclusion of all others.” This is completely consistent with our approach, which characterizes the *objective of the econometrician* in terms of a scalar (the log-likelihood of a mixture model) derived

This approach assumes that any one observation can be generated by both criteria, although it admits of extremes in which one or other criteria wholly generates the observation. One could alternatively define a grand likelihood in which observations or subjects are classified as following one criteria or the other on the basis of the latent probabilities π^{SP} and π^{A} . El-Gamal and Grether [1995] illustrate this approach in the context of identifying behavioral strategies in Bayesian updating experiments. In the case of the SP/A model, it is natural to view the tension between the criteria as reflecting the decisions of a given individual for a given task. Thus we do not believe it would be consistent with the SP/A model to categorize *choices* as wholly driven either of SP or A.

These priors also imply that we prefer not to use mixture specifications in which *subjects* are categorized as completely SP or A. It is possible to rewrite the grand likelihood (10) such that $\pi_i^{\text{SP}} = 1$ and $\pi_i^{\text{A}} = 0$ if $I_i^{\text{SP}} > I_i^{\text{A}}$, and $\pi_i^{\text{SP}} = 0$ and $\pi_i^{\text{A}} = 1$ if $I_i^{\text{SP}} < I_i^{\text{A}}$, where the subscript i now refers to the individual *subject*. The general problem with this specification is that it assumes that there is no effect on the probability of SP and A from task domain. We do not want to impose that assumption, even for a relatively homogenous task design such as ours.

4. Results

Table 1 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 1 displays the implied probability weighting function and decision weights. The decision weights are 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

from modeling the *objective of the latent decision-maker* in terms of a probability distribution defined over two or more criteria.

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 0.302 and significantly greater than zero. Thus we infer that there is some evidence of concave utility as well as probability weighting. Constraining the utility function to be linear in the RDEV specification slightly increases the curvature of the probability weighting function, as one might expect.

Table 2 and Figure 2 show the results of estimating the SP/A model. First, we find evidence that the utility function is concave, since $\rho > 0$ and has a 95% confidence interval between 0.59 and 0.69. Hence it would be inappropriate to assume RDEV for the SP part of the SP/A model in this high-stakes domain, exactly as Lopes and Oden [1999; p.290] conjecture.

Second, we find evidence that the SP weighting function is initially concave and then convex in probabilities. In the jargon of the SP psychological processes underlying this weighting function, this indicates that potential-minded attitudes dominate the security-minded attitudes for smaller probabilities, but that this is reversed for higher probabilities. In the *DOND* context, the probabilities are symmetric, and significantly less than $\frac{1}{2}$ for virtually all rounds. Hence the predominant attitude implied by $\gamma < 1$ is that of potential-minded attitudes.

Third, the estimates of χ , ξ and ψ for the aspiration weighting function imply that it is steadily increasing in the prize level, with the concave shape shown in Figure 2. At a prize level of roughly £44,000 the aspiration threshold is $\frac{1}{2}$. This function does not assign zero weight to prizes below that level, although the functional form effectively allowed that. If we round prizes to the nearest £10,000, the average aspiration weights are 0.09, 0.22, 0.34 and 0.46 for each of the prizes from £0 up to £40,000. If £44,000 seems optimistic, and it is given the historical evidence, recall that this is just one of two decision criteria that the contestant is assumed to use in making *DOND*

decisions. The other criterion combines security and potential considerations, as noted above.

Finally, the two component processes of the SP/A model each receive significant weight overall. We estimate that the weight on the SP component, π^{SP} , is 0.36, with a 95% confidence interval between 0.31 and 0.41. A formal hypothesis test that the two components receive equal weight can be rejected at a p -value of less than 0.001, but each component clearly plays a role in decision-making in this domain.

5. Model Comparisons

Mixture models provide a natural way to compare the predictive power of alternative models of behavior, whether or not the models are nested.¹⁸ The estimates for the role of the SP and A criteria within the SP/A model already indicate that decision makers appear to use both the familiar “risk-utility” tradeoffs of traditional economics models, as well as the notion of an income threshold.

To the extent that the SP component is simply a re-statement of the RDU model, this SP/A model nests RDU, so this result provides some evidence in favor of the SP/A notion that one needs two criteria to appropriately characterize behavior in *DOND*. Furthermore, since the RDU model nests at least one parametric version of EUT, these results indicate that behavior is inconsistent with that version of EUT.

We can extend our analysis to consider the possibility that some choices or decision makers are consistent with EUT and some are consistent with RDU. In effect, this hypothesis suggests a nested mixture model: at the top level one allows EUT and SP/A latent decision-making process explain behavior, and at the bottom level *within* the SP/A process one allows latent SP and A processes to explain behavior. This hypothesis is not the same as the hypothesis that the SP criterion (which is the RDU criterion) collapses to EUT. In effect, this hypothesis is that there are *three* latent decision-making processes at work: EUT, RDEU and the A criterion. Assuming a CRRA utility

¹⁸ Harrison and Rutström [2005] point out that the older statistical literature on non-nested hypothesis tests evolved as a “second best” alternative to being able to estimate finite mixture models.

function, we estimate that the EUT process accounts for only 5.4% of observed choices, with the relative weights on the SP/A criterion accounting for the rest. In turn, the relative weights on the SP and A criteria remain essentially the same as when we assumed that none of the choices were generated by EUT decision-makers.

Thus we conclude that the weight of the evidence supports the SP/A model over the assumption that behavior is characterized *solely* by EUT or even by RDU decision making.

6. Conclusions

We provide a formal statement and application of a model of choice under uncertainty from psychology that has been neglected by economists, but which has many interesting features. First, and foremost, it explicitly employs multiple criteria for the evaluation of prospects. This characteristic is intuitively appealing, and informally used in many expositions of alternatives to EUT; examples include decision weights, loss aversion, and regret or disappointment aversion. In many cases these criteria can be conveniently collapsed to a single criterion, but in some cases they cannot; examples include appeals to income thresholds, editing processes, the use of similarity relations, and other heuristics from psychology.¹⁹

We demonstrate how one can obtain full maximum likelihood estimates of the SP/A model, and integrate the dual decision criteria in a natural decision-theoretic and statistical manner. These methodological insights extend to applications of the SP/A model to other settings, as well as to other multiple-criteria decision models.

We apply the model to a rich domain in which prizes are highly skewed, and where it is plausible to expect individuals to have income thresholds that might affect behavior in addition to familiar utility-risk tradeoffs. Our statistical results allow the data to determine the relative weight of the two criteria, and do not *a priori* constrain behavior to use either or both. We find evidence that

¹⁹ See Starmer [2000] on editing processes, Tversky [1977], Luce [1956], Rubinstein [1988] and Leland [1994] on similarity relations, and Brandstätter, Gigerenzer and Hertwig [2006] on the myriad of heuristics proposed in the broader psychology literature.

both criteria play a role in explaining behavior. Although the specific weights attached to each criteria might be expected to vary from task domain to domain, we find that nearly two-thirds of the weight is on the aspiration criterion.

Table 1: Estimates for *Deal or No Deal* Game Show Assuming RDU

Parameter	Estimate	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
A. RDEV, assuming utility is linear in prizes				
γ	0.342	0.011	0.322	0.363
μ	0.060	0.004	0.053	0.068
B. RDU				
ρ	0.302	0.014	0.274	0.330
γ	0.550	0.022	0.506	0.595
μ	0.085	0.004	0.077	0.092





Table 2: Estimates for *Deal or No Deal* Game Show Assuming SP/A

Parameter	Estimate	Standard Error	Lower 95% Confidence Interval	Upper 95% Confidence Interval
ρ	0.654	0.027	0.601	0.707
γ	0.664	0.052	0.562	0.767
χ	1.355	0.146	1.068	1.642
ξ	5.330	1.510	2.371	8.289
ψ	0.002	0.001	-0.001	0.005
μ	0.183	0.029	0.125	0.241
π^{SP}	0.353	0.025	0.304	0.401
$\pi^{\text{A}}=1-\pi^{\text{SP}}$	0.647	0.025	0.599	0.696



References

- 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, "Dynamic Choice Behavior in a Natural Experiment," *Working Paper 06-10*, Department of Economics, College of Business Administration, University of Central Florida, 2006.
- Andersen, Steffen; Harrison, Glenn W., Lau, Morten I., and Rutström, E. Elisabet, "Risk Aversion in Game Shows," in J.C. Cox and G.W. Harrison (eds.), *Risk Aversion in Experiments* (Greenwich, CT: JAI Press, Research in Experimental Economics, Volume 12, 2007 forthcoming).
- Ballinger, T. Parker, and Wilcox, Nathaniel T., "Decisions, Error and Heterogeneity," *Economic Journal*, 107, July 1997, 1090-1105.
- Baumol, William J., "An Expected Gain-Confidence Limit Criterion for Portfolio Selection," *Management Science*, 10, 1963, 174-182.
- Bell, David E., "Risk, Return, and Utility," *Management Science*, 41(1), January 1995, 23-30.
- Brandstätter, Eduard; Gigerenzer, Gerd, and Hertwig, Ralph, "The Priority Heuristic: Making Choices Without Trade-Offs," *Psychological Review*, 113(2), 2006, 409-432.
- Byrne, R.; Charnes, A.; Cooper, W.W., and Kotakek, K., "A Chance-Constrained Approach to Capital Budgeting with Portfolio Type Payback and Liquidity Constraints and Horizon Posture Controls," *Journal of Financial and Quantitative Analysis*, 2(4), December 1967, 339-364.
- Byrne, R.; Charnes, A.; Cooper, W.W., and Kotakek, K., "Some New Approaches to Risk," *The Accounting Review*, 43(1), January 1968, 18-37.
- Camerer, Colin; Babcock, Linda; Loewenstein, George, and Thaler, Richard, "Labor Supply of New York City Cabdrivers: One Day at a Time," *Quarterly Journal of Economics*, 112, May 1997, 407-441.
- Campbell, John Y.; Lo, Andrew W., and MacKinlay, A. Craig, *The Econometrics of Financial Markets* (Princeton: Princeton University Press, 1997).
- 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.
- El-Gamal, Mahmoud A., and Grether, David M., "Are People Bayesian? Uncovering Behavioral Strategies," *Journal of the American Statistical Association*, 90, 1995, 1137-1145.
- Farber, Henry S., "Is Tomorrow Another Day? The Labor Supply of New York City Cabdrivers," *Journal of Political Economy*, 113(1), 2005, 46-82.

- Fishburn, Peter C., "Additive Utilities with Incomplete Product Sets: Application to Priorities and Assignments," *Operations Research*, 15(3), May-June 1967, 537-542.
- 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., 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.
- Hey, John D., "Experimental Investigations of Errors in Decision Making Under Risk," *European Economic Review*, 39, 1995, 633-640.
- Hey, John D., and Orme, Chris, "Investigating Generalizations of Expected Utility Theory Using Experimental Data," *Econometrica*, 62(6), November 1994, 1291-1326.
- Hogarth, Robin M., *Educating Intuition* (Chicago: University of Chicago Press, 2001).
- Johnson, Norman L.; Kotz, Samuel, and Balakrishnan, N., *Continuous Univariate Distributions, Volume 2* (New York: Wiley, Second Edition, 1995).
- Kahneman, Daniel, and Tversky, Amos, "Prospect Theory: An Analysis of Decision Under Risk," *Econometrica*, 47, 1979, 263-291.
- Keeney, Ralph L., and Raiffa, Howard, *Decisions with Multiple Objectives: Preferences and Value Tradeoffs* (New York: Wiley, 1976).
- Kirkwood, Craig W., *Strategic Decision Making: Multiobjective Decision Analysis with Spreadsheets* (Belmont, CA: Duxbury Press, 1997).
- Kleinmütz, Benjamin, "Why we still use our heads instead of formulas: Toward an integrative approach," *Psychological Bulletin*, 107, 1990, 296-310; reprinted in T. Connolly, H.R. Arkes & K.R. Hammond (eds.), *Judgement and Decision Making: An Interdisciplinary Reader* (New York: Cambridge University Press, 2000).
- Leland, W. Jonathan, "Generalized Similarity Judgements: An Alternative Explanation for Choice Anomalies," *Journal of Risk & Uncertainty*, 9, 1994, 151-172.
- Liang, K-Y., and Zeger, S.L., "Longitudinal Data Analysis Using Generalized Linear Models," *Biometrika*, 73, 1986, 13-22.
- Liberatore, Matthew, and Nydick, Robert, *Decision Technology: Modeling, Software, and Applications* (New York: Wiley, 2002).
- Loomes, Graham; Moffatt, Peter G., and Sugden, Robert, "A Microeconomic 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.
- Lopes, Lola L., "Algebra and Process in the Modeling of Risky Choice," in J.R. Busemeyer, R. Hastie & D.L. Medin (eds), *Decision Making from a Cognitive Perspective* (San Diego: Academic Press, 1995).
- Lopes, Lola L., and Oden, Gregg C., "The Role of Aspiration Level in Risky Choice: A Comparison of Cumulative Prospect Theory and SP/A Theory," *Journal of Mathematical Psychology*, 43, 1999, 286-313.
- Luce, R. Duncan, "Semiororders 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.
- Oden, Gregg C., and Lopes, Lola L., "Risky Choice With Fuzzy Criteria," *Psychologische Beiträge*, 39, 1997, 56-82.
- Quiggin, John, "A Theory of Anticipated Utility," *Journal of Economic Behavior & Organization*, 3(4), 1982, 323-343.
- Rogers, W. H., "Regression standard errors in clustered samples," *Stata Technical Bulletin*, 13, 1993, 19-23.
- Roy, A.D., "Safety First and the Holding of Assets," *Econometrica*, 20(3), July 1952, 431-449.
- Roy, A.D., "Risk and Rank or Safety First Generalised," *Economica*, 23, August 1956, 214-228.
- 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.
- Saaty, Thomas L., *The Analytic Hierarchy Process* (New York: McGraw Hill, 1980).
- Starmer, Chris, "Developments in Non-Expected Utility Theory: The Hunt for a Descriptive Theory of Choice Under Risk," *Journal of Economic Literature*, 38, June 2000, 332-382.
- 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.
- 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, "Intransitivity of Preferences," *Psychological Review*, 76, 1969, 31-48.
- Tversky, Amos, "Features of Similarity," *Psychological Review*, 84, 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).
- von Winterfeldt, Detlof, and Edwards, Ward, *Decision Analysis and Behavioral Research* (New York: Cambridge University Press, 1986).
- 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(2), May 2003, 133-138.
- Yaari, Menahem E., "The Dual Theory of Choice under Risk," *Econometrica*, 55(1), 1987, 95-115.

Appendix: Estimation Procedure (NOT FOR PUBLICATION)

The basic logic of our approach can be explained from the data and simulations shown in Table A1, and assuming for simplicity that the decision maker behaves as if using EUT to evaluate choices. We discuss the non-EUT specification in the text, but the basic estimation logic is the same. Complete details are provided in Andersen, Harrison, Lau and Rutström [2007].

There are 6 rounds in which the banker makes an offer, and in round 7 the surviving contestant simply opens his box. In the tabulations shown in Table A1 we observed 461 contestants play the game. Only 65, or 15%, 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 £10,184 with a standard deviation of £9,575. 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-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 £12,532 with a standard deviation of £12,107. Similarly for each of the other virtual lotteries shown.

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 A1. 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 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 A1 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 100,000 simulations for *each* player in *each* round, so as to condition on their history.²⁰ The

²⁰ If bank offers were a deterministic and known function of the expected value of unopened prizes, we would not need anything like 100,000 simulations for later rounds. For the last few rounds of a full game,

fraction of the EV of unopened prizes is also round-specific, and is a draw from a normal distribution for that round based on observed data. Thus the simulated offer received in any future, virtual round is uncertain, due to uncertainty in the EV of unopened cases in that future round *and* uncertainty in the fraction of that EV that the banker will use.²¹

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.²² 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 $v(1)=6, v(2)=5, \dots, v(6)=1$ as shown in Table A1. 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.²³

To state the estimation problem more formally, assume that utility is defined over money m using a Constant Relative Risk Aversion (CRRA) function

$$u(m) = m^{1-r} / (1-r) \tag{A1}$$

where $r \neq 1$ 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 function has been popular in the literature, since it requires only one parameter to be estimated.

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

$$EU_i = \sum_{k=1, 20} [p_k \times u_k]. \tag{A2}$$

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

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

²¹ The simulated EV is subject-specific and round-specific; the fraction of the EV used by the banker for that is not subject-specific, but is round-specific.

²² 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.]).

²³ The only complication from using a 20-point approximation might occur when one undertakes probability weighting. However, if one uses rank-dependent probability weighting this issue disappears. For example, a 4-point virtual lottery with prizes 100, 100, 200 and 200, each occurring with probability $1/4$, is the same as a lottery with prizes 100 and 200 each occurring with probability $1/2$.

$$\nabla EU = (EU_{BO} - EU_L) / \mu \quad (A3)$$

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].²⁴ The index ∇EU is then used to define the cumulative probability of the observed choice to “Deal” using the cumulative standard normal distribution function:

$$G(\nabla EU) = \Phi(\nabla EU). \quad (A4)$$

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

$$\ln L^{EUT}(r, \mu; y) = \sum_i [(\ln G(\nabla EU) \mid y_i=1) + (\ln (1-G(\nabla EU)) \mid y_i=0)] \quad (A5)$$

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 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 boxes, 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 $\Gamma=100,000$). 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

²⁴ 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], 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.

the expected offer in round 5. And for a virtual lottery that only involves looking forward to rounds 4, 3 and 2, respectively. Table A1 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 A1.

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 A1. 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.²⁵ In fact, other versions of *DOND* can be used to illustrate the flexibility of this approach, since they sometimes employ “follow on” games that can simply be folded into the virtual lottery simulation.²⁶ 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 in our application to the SP/A model and in Andersen, Harrison, Lau and Rutström [2006].

The one weakness of this approach is that it only approximates the fully dynamic solution of

²⁵ 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.

²⁶ For example, the *Chance* and *Supercase* variants in the Australian version.

the problem facing the contestant. To illustrate, consider the contestant in round 6, facing 5 unopened prizes and having to open 3 prizes if he declines the bank offer in round 6. Call these prizes V, W, X, Y and Z. There are 9 combinations of prizes that could remain after opening 3 prizes. Our approach to the VL, from the perspective of the round 6 decision, evaluates the payoffs that confront the contestant for each of these 9 combinations if he “mentally locks himself into saying D in round 7 and then gets the stochastic offer given the unopened prizes” or if he “mentally locks himself into saying ND in round 7 and then opens 1 more prize.” The former is the VL associated with the strategy of saying ND in round 6 and D in round 7, and the latter is the VL associated with the strategy of saying ND in round 6 and ND again in round 7. We compare the EU of these two VL as seen from round 6, and pick the largest as representing the EU from saying ND in round 6. Finally, we compare this EU to the U from saying D in round 6, since the offer in round 6 is known and deterministic. The simplification comes from the fact that we do not evaluate the utility function in the each of the possible *virtual* round 7 decisions. A complete enumeration of each possible path would undertake 9 paired comparisons.

As one applies backwards induction to this complete enumeration, the number of possible paths explodes, and becomes computationally tedious to evaluate. Andersen, Harrison, Lau and Rutström [2007; Appendix B] demonstrate that the approximation provided by our approach is an excellent one. The main reason it is a good approximation is that the critical feature of the forward-looking path in *DOND* is to account for the expected increase in bank offers as a percentage of the expected value of unopened prizes, and our approach allows for that.

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.²⁷

²⁷ 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 individual heterogeneity, and show that it is statistically significant in their models (which ignore dynamic features of the game).

Table A1: Virtual Lotteries for British *Deal or No Deal* Game Show

Average (standard deviation) of payoff from virtual lottery using 100,000 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	461 100%	0	£6,157	£10,184 (£9,575)	£12,532 (£12,107)	£13,809 (£14,126)	£15,772 (£19,444)	£25,772 (£37,820)	£25,773 (£55,916)
2	461 100%	0	£9,265		£12,321 (£11,961)	£13,580 (£13,966)	£15,508 (£19,236)	£25,351 (£37,457)	£25,346 (£55,200)
3	461 100%	23	£11,527			£13,568 (£13,913)	£15,496 (£19,182)	£25,320 (£37,208)	£25,321 (£54,877)
4	438 95%	99	£12,854				£15,496 (£19,226)	£25,328 (£37,237)	£25,30 (£54,842)
5	339 75%	152	£14,190					£23,842 (35,575)	£23,833 (£52,430)
6	187 39%	65	£13,704						£19,584 (£45,912)
7	122 26%								

Note: Data drawn from observations of 461 contestants on the British game show, plus author's simulations of virtual lotteries as explained in the text.