Risk Aversion and Expected Utility Theory: An Experiment with Large and Small Stakes.

Matilde Bombardini and Francesco Trebbi*

First Draft November 2005
This Draft August 2010

Abstract

We employ a novel data set to estimate a structural econometric model of the decisions under risk of players in a game show where lotteries present payoffs in excess of half a million dollars. Differently from previous studies in the literature, the decisions under risk of the players in presence of large payoffs allow to estimate the parameters of the curvature of the vN-M utility function not only locally but also globally. Our estimates of relative risk aversion indicate that a constant relative risk aversion parameter of about one captures the average of the sample population. In addition we find that individuals are practically risk neutral at small stakes and risk averse at large stakes, a necessary condition, according to Rabin (2000) calibration theorem, for expected utility to provide a unified account of individuals' attitude towards risk. Finally, we show that for lotteries characterized by substantial stakes non-expected utility theories fit the data equally well as expected utility theory.

JEL Classification codes: D81, C99

We would like to thank Endemol Italia S.p.A for kindly providing access to the recorded episodes of the show "Affari Tuoi". In particular we would like to thank Luca Passerini and Salvatore Spirlí for extensive discussion about the game rules and the selection process of the participants and Silvia Brasca for facilitating the data collection. We thank Paul Beaudry, Gorkem Celik, David Green, Jerry Green, Drew Fudenberg, Yoram Halevy, David Laibson, Thomas Lemieux, and Okan Yilankaya for useful comments and suggestions. Particularly, Philippe Aghion, Alberto Alesina, and Andrei Shleifer are gratefully acknowledged for numerous discussions. Kareem Carr provided excellent research assistance. We are also grateful to seminar participants at the University of British Columbia, the 12th conference on Foundations and Applications of Utility, Risk and Decision Theory, and Harvard University.

*University of British Columbia, Department of Economics.

1 Introduction

In its most basic formulation decision theory offers a simple way to compare the degree of risk aversion across individuals. Individual A is globally more risk averse than individual B if the certainty equivalent of lottery L for individual A is lower than the certainty equivalent for individual B and if this is true for any lottery L. Armed with such a simple method, the ideal experiment to measure risk aversion would be to present subjects with simple lotteries and ask them to reveal their certainty equivalent to each lottery. Under specific parametric assumptions about the utility function it is possible to translate these certainty equivalents into a quantity that represents the degree of risk aversion of the individual. Further, according to expected utility theory, risk aversion derives from the curvature of the utility of money, so such experiment would require to vary the stakes of the lotteries proposed in order to trace out the shape of the utility of money.

What we just described is very similar to the setup involved in the Italian game show "Affari Tuoi" which is the subject of this paper. A detailed description of the game rules is provided in the following section, but the main choice contestants make in this game can be easily summarized in the following example. At a given stage of the game the contestant is holding a closed box that might contain either 1 euro or $100,000$ euros. At this point the contestant is offered $33,000$ euros and must decide whether to accept this certain payoff or face a lottery where the prizes 1 and $100,000$ euros are equally likely. If the contestant accepts the offer, then $33,000$ euros is an upper bound on her certainty equivalent. If the contestant rejects the offer, then $33,000$ euros is a lower bound on her certainty equivalent[1]. A given contestant is asked to make this simple choice several times and facing different lotteries during the same episode, therefore we observe multiple decisions for the same individual. From such decisions one is able to estimate risk preferences, the first objective of this paper.

According to Harrison and List (2004) what we study in this paper is a 'natural experiment' since we did not intervene in designing the rules of the game. Nevertheless it presents several features that we would have chosen, were we to design such an experiment. First, observe that in order to mimic the ideal experiment that we described above, one has to design the experiment so that the subject is asked to reveal the certainty equivalent of the lottery. Truthful revelation is possible to enforce only with incentive compatible schemes (BDM for instance). In this game the fact that the offer is take-or-leave and is made only once to the same person prevents the subject from playing strategically and guarantees that the choice involves only risk preferences. Second, given the established anomalies in the behavior of individuals in the presence of infinitesimally small probabilities placed on extremely large payoffs, such as in the case of lottery tickets, it is desirable to present the subject with few prizes and finite and identical probabilities. This also guarantees that the choice is not affected by the ability to assign the correct probabilities to each outcome. Third, as required to recover the entire shape of the utility of money, the experiment entails small stakes lotteries, involving 1 or 500 euros as well as large stakes lotteries, where the payoff can reach $500,000$ euros.

We report three main findings, with implications that span beyond decision theory. First, we find that,

[1] This game was actually played by Maria Grazia Rovai, a housekeeper from Vasto, on May 17, 2005. She rejected the offer and won 1 euro.

on average, logarithmic utility approximates well preferences of contestants. Confirmation that log utility is a good description of preferences over lotteries with large stakes is extremely valuable, given the very sizable portion of papers in the macroeconomics and finance literature assuming log utility.

However, and this is our second contribution, we also detect substantial heterogeneity in relative risk aversion parameters, with standard deviations around 1. The role of second moments is important for those researchers aiming at providing micro-foundations of aggregate variables, like the aggregate savings rate or a country's aggregate portfolio allocations (especially when aggregation is not trivial due to non-linearities). Here we offer a benchmark figure for future calibrations of the dispersion of risk tolerance[2].

Our third contribution is in showing that Expected Utility Theory (EUT) is a good description of individual behavior under risk. Not only the number of contestants choice which cannot be rationalized by EUT within our estimation framework is minimal, but also we provide evidence of local risk neutrality.

Indeed, the ability to trace out the utility of money over what we believe is a relevant range allows us to address the point raised by Rabin (2000) who predicts that "data sets dominated by modest-risk investment opportunities are likely to yield much higher estimates of risk aversion than data sets dominated by larger-scale investment opportunities". Rabin implies that the same underlying utility of money cannot be used when considering small and large-stake lotteries. We find that logarithmic utility fits the data accurately when we consider both small and large-stake lotteries. In Figure 1A the horizontal axis measures the expected value of the lottery (in our example above it corresponds to $50,000$ euros) while the vertical axis measures the offer made to a contestant that faced such lottery. Each lottery is a simple two-payoff object, just like our introductory example. In general, accepted offers lie above points indicating rejected offers for a given expected value of the lottery. Moreover if we locally fit a line through the offers corresponding to each lottery (therefore tracing the locus of indifference, the solid line) we find a curve that almost coincides with the logarithmic function represented on the same graph. For comparison, notice that the linear utility and a CRRA utility with a coefficient of risk aversion equal to 2 do a much poorer job at describing the data. In Figure 1B, where lotteries have a five-payoffs structure, the story is similar.

Consider now what the Rabin's calibration theorem predicts. If we estimate anything but almost risk neutrality at small stakes then we are bound to observe unrealistically high risk aversion at large stakes. This time a graphical representation helps showing how close our results come to Rabin's prediction. Figure 2 shows the subset of lotteries shown in Figure 1A with prizes below $18,000$ euros. As Rabin predicts, at lower stakes the subjects are almost risk neutral, as the comparison between logarithmic utility and linear utility clearly shows. The message emerging from these graphs and from the quantitative results presented

[2]Indeed, there is increasing interest in modeling and calibrating risk tolerance heterogeneity across several fields within economics. Schulhofer-Wohl (2008) studies heterogeneous risk preferences and the welfare cost of business cycle fluctuations. In health economics Cutler, Filkenstein, and McGarry (2008) discuss how heterogeneity in risk preferences may be more relevant for explaining insurance demand than heterogeneity in risk. Heterogeneity in risk preferences has been attracting attention also in portfolio theory, where the goal is to explain the massive variation in cross-sectional portfolio composition. For instance, a majority of households does not own stocks and a sizeable group of stockholders mostly owns stock from the firm they own or they work for (and appear hence underdiversified). In development economics Mazzocco (2004) and Mazzocco and Saini (2008) focus on heterogeneous risk preferences and discuss its (not obvious) implications for efficient risk sharing of households.

in the paper is that we do not need to distinguish between expected utility of income versus expected utility of wealth in order to free EUT of Rabin's criticisms, as argued by Cox and Sadiraj (2006). As shown later in the paper, both reference points (wealth and income) offer qualitatively the same conclusions. Therefore not only our results present a useful benchmark in measuring risk preferences when the individual decisions are important and involve significant amounts of money, which to the best of our knowledge is a new result, but also produce evidence in favor of EUT.

1.1 Related literature

In light of all the desirable features of this experiment, we see many of the previous experiments set out to assess risk preferences to lack one or more such ideal conditions.

As pointed out by Harrison and List (2004) many of the experiments conducted in the lab involve college students because of availability and proximity to academic researchers. In order to obtain more representative subject pools, field experiments draw samples on populations with wider demographics. Similarly, the subjects of this experiment differ in terms of age, education, income, geographical origins in a way that, as we will later explain, makes them reasonably representative of the population. This is not the case for other game show studies like Metrick (1995) where the author admits the contestants are "wealthier and better educated than the average person".

The issue of ability is not just a concern from a selection or self-selection point of view, it is possibly a more important concern if we believe it cannot be separated from risk preferences. This might happen if ability, like in the case of 'Jeopardy!', affects the assessment of one's probability to answer a question or the calculation of fairly complicated winning odds like in 'Card Sharks' and 'Lingo'. Using the former game show, Metrick (1995) is constrained to employing the average probability that the contestant will answer a question correctly to assess the degree of risk aversion of the participant by looking at how much the subject wagers. It is clear that the participant's own assessment of her ability to answer a future question might interfere with risk preferences in determining the amount wagered. Metrick's results do not allow him to reject the hypothesis of risk neutrality[3].

In his study of the game show 'Card Sharks' Gertner (1993) considers contestants who calculate the probability of picking, given a certain card, a higher-ordered card from a deck without replacement. Beetsma and Schotman (2001) consider the game show 'Lingo', where contestants have to evaluate the probability of extracting balls from an urn without replacement. In these cases the ability to compute probabilities might interfere with risk preferences and it could be the reason why Gertner (1993) and Beetsma and Schotman (2001) obtain relatively high coefficients of relative risk aversion.[4]

A number of studies have made use of data from various editions of the same game we analyze here. To the best of our knowledge this was the first study to propose a fully structural estimation of this natural

[3] When probabilities are subjective though, over-confidence can be mistaken for risk neutrality, as Metrick himself recognizes.

[4] Gertner's point estimate is 4.8 which is considerably higher than what we obtain in this study. Beetsma and Schotman's estimates vary depending on the reference income level, but are well above 3 for income levels comparable to the ones used in our study.

experiment. A very detailed and clear summary of the methods employed and the results obtained is provided by Andersen, Harrison, Lau and Rutstrom (2008), who also propose a general estimation strategy for this class of games. These subsequent studies vary in the type of alternative decision theoretic approaches they analyze and the methodology employed. Andersen et al. (2006b), DeRoos and Sarafidis (2009), Mulino, Scheelings, Brooks and Faff (2009) follow our structural approach of estimating the full dynamic choice model by maximum likelihood, while Post, Van den Assem, Baltussen and Thaler (2008), Blavatskyy and Pogrebna (2010), Deck, Lee and Reyes (2006) use a reduced-form approach.

Compared also to non-game settings, the experiment in this study exhibits considerable advantages. Take for example the work of Jullien and Salanié (2000), where racetrack bets are used to estimate risk preferences. They are constrained to adopting a representative agent approach due to the lack of information on individual bettors. In our sample we are able to identify a few important individual characteristics and therefore allow for heterogeneity in risk preferences. Most important among these is income, which allows us to adopt constant relative risk aversion utility. Allowing for unobserved heterogeneity has appealing features as shown by Cohen and Einav (2007). In their paper they present an interesting application of this methodology, albeit in a more complex environment (the choice of insurance deductibles) that requires accounting for adverse selection and moral hazard, which substantially complicates the estimation.

This and other studies, making use of insurance decisions to evaluate risk preferences like Cicchetti and Dubin (1994), involve more complicated lotteries than the ones we analyze in this paper and payoffs that are several orders of magnitude smaller. The nature of the lottery might be problematic since individuals are expected to make a decision about buying insurance based on knowledge of the exact probability of a car accident or telephone line failure. Arguably a 50-50 lottery over two cash prizes is a simpler problem for the individual to evaluate. The size of the lottery stakes is also problematic for these studies because insurance deductibles and telephone line failures do not involve sums larger than a few hundred dollars. Such small bets simply do not allow to trace out significant portions of the utility of income.

The paper is organized as follows. Section 2 describes the data and presents the principal features of the game show in detail. Section 3 presents a model of the contestant's dynamic decision problem and its solution. Section 4 reports the estimation results of the model and offers an empirical interpretation of Rabin (2000). Section 5 extends the econometric analysis to non-expected utility paradigms and provides a comparison with the estimates obtained under expected utility theory. Section 6 presents our conclusions.

2 The Game

Structure

In an episode of "Affari Tuoi" (literally "Your Business") each of 20 participants, representative of each Italian region, is assigned a box, numbered from 1 to 20[5]). Each box contains a prize which can be monetary or in-kind[6]. Winnings are always paid through a wire transfer so objects represent a certain monetary

[5] The assignment of the box to the region is made in secret by a notary.

[6] The notary randomly assigns each prize to each box. In-kind prizes are assigned a monetary value by RAI (Radiotelevisione

amount and serve the only purpose of entertainment[7]. Table 1A reports the distribution of prizes contained in the twenty boxes. It is worth noticing that these prizes are considerably large, not only compared to the usual experimental settings[8], but also relative to previous studies of game shows. For instance Gertner (1993) employs results from the game 'Card Sharks' where the stakes are up to $\$16,000$ and the average winnings are $\$3,200$. In Metrick (1995) the potential stakes are higher but winnings above $\$50,000$ represent only 6 percent of the total. In this game the set of prizes include $100,000$ euros, $250,000$ euros, up to $500,000$ euros. The stakes are sizable compared to lifetime wealth for some of the participants.

The game consists of two phases: an introductory part and the main game. The first part consists of a single question. The person who gets the right answer[9] in the shortest time gains access to the main phase of the game. This phase is the subject of our study. In this part of the game the contestant is interviewed shortly by the host and usually reveals some information about herself that we employ in this study. We do not have access to any information beyond the details provided by the participants themselves[10]. We always know the region of provenience of the participant and the gender. Usually the participant reveals whether she is from a specific city or from the countryside, her occupation, whether she is married and has any children. Less often participants state their age. Where the exact age is missing, we produced a range to describe the approximate age of the participant, generally a 10-year window.

The main game begins with the participant holding the originally assigned box. She is asked to pick one of the remaining nineteen boxes. The chosen box is opened and the prize contained is revealed. The opened box and its corresponding prize drop out of the set of possible winnings for the participant. In the same fashion the participant chooses and opens five more boxes. Every time a box is opened it drops out of the game along with the prize it contains. After the first six boxes are opened an offer is made. An offer can take the form of a monetary amount or of a "change". If the offer is a monetary amount, the participant has the option of accepting the amount of money offered and abandon the game or to continue and open more boxes. When the offer is a change the participant has the option of substituting the box she holds in her hands with one of the remaining unopened boxes. If the monetary offer is rejected or if the offer is in the form of a change then 3 more boxes are opened before the next offer, and so on. If no offer is ever accepted the contestant wins the prize in the box she is holding. This implies that there are up to a maximum of five offers per episode[11]. At the final (fifth) offer in particular there are two boxes unopened, one of which is in the hands of the participant. This is, along with the fourth offer, when most of the 'drama' of the game unravels as the participant often faces very different potential prizes and is torn between accepting a safe offer, which is always lower than the mean of the two prizes, and facing a lottery. Table 1B shows the timing

Italiana) and they never exceed 500 euros value.

[7] All prizes with face value smaller than 50 euro are paid 50 euros. This is accounted for in the estimation.

[8] For experiments involving relatively high monetary incentives see Holt and Laury (2002) and the study on Beijing University participants in Kachelmeier and Shehata (1992).

[9] The question is usually of very peculiar nature and does not appear to reflect contestant's ability or knowledge more than sheer luck in guessing. An example would be "How many churches in Italy are dedicated to Saint Paul?". Random access to the main section of the game does not seem an unreasonable assumption in this light.

[10] The Appendix contains the details concerning variables' definition and construction.

[11] There are only few exceptions to this pattern. In some episodes offers are made more frequently i.e. after one box is opened.

of the game.

The offers are made by a person, named by the host as the “Infame”, literally the “Infamous”, who communicates with the host through a telephone during the show. The offerer knows the content of the box in the hands of the participant[12] and decides the offer to make based solely on the behavior of the participant and the information that is revealed about the player during the episode[13]. When we interviewed the “Infame” he admitted not using any specific algorithm or any aid other than his own experience. What we observe in the data is that there is a positive correlation between the content of the box held by the contestant at the time the offer is made and the amount of money offered, controlling for the size of the lottery. For instance the correlation between the monetary offer (relative to the expected value of the lottery) and the relative content of the box held by the contestant is 0.36 at the fifth offer and 0.32 at the fourth offer[14]. This positive correlation suggests that the offerer finds it profitable to “convince” the contestant to forego a high prize by offering a relatively high amount more frequently than when the contestant has nothing in her hands[15]. In an online Appendix we suggest a theoretical framework to describe the equilibrium behavior of the offerer. In the empirical estimation we take the behavior of the offerer as given and use the equilibrium strategy that we observe[16].

Selection

The show “Affari Tuoi” which means literally ‘your business’ aired for the first time on the Italian television channel RAI 1 on October, 13th 2003 and has been on the air for four seasons[17]. The show was interrupted for the summer on June, 3rd 2005 and resumed in the fall of 2005 with a new host[18]. Endemol, the producing company and a Dutch multinational, is the owner of the show’s format and holds the rights to produce and lease the format to the various TV stations. The show is vastly popular in Italy with peaks of

[12] We would like to emphasize the point that the “banker”, or offerer, knows the content of the boxes. We personally spoke to the producers of the game and to the Infamous himself. Notice that this is known by the contestants themselves and it emerges clearly by watching even a few episodes of the game. Indeed, the nickname Infamous derives from the fact that the “banker” tries to mislead the contestants.

[13] The host himself is not aware of the content of the box held by the contestant.

[14] This correlation is robust at least at the 5% confidence level to the inclusion of controls such as the standard deviation of prizes and the square of the relative content of the box held by the contestant at both the fourth and the fifth offer. Regression results documenting this robust relationship are available from the authors.

[15] An anonymous referee rightly wonders whether, given the amount of money involved in advertising, the TV station even cares about the amount of money disbursed or it is just trying to improve the entertainment value of the show. This is a legitimate question since we argue that the Infamous is playing strategically and minimizing the amount of money paid to contestants. We have three arguments in support of a cost-minimizing Infamous: (i) we interviewed one of the producers and he confirmed that, although the goal was not to minimize the payout, the show did have a budget per episode of 30,000 euros (which the producers managed to maintain) and this was a long-run constraint which obviously did not bind episode by episode; (ii) we interviewed the Infamous, who admitted trying to “trick” people and convince them to accept the offer when they had the high prize in their hands (without making excessive offers that would reveal the content of the box); (iii) by watching many episodes we understood that the entertainment value of the show was actually increased by the attempt, by the Infamous, to reduce the amount of money disbursed.

[16] Further discussion of the advantages of such approach is provided in the following section.

[17] The four seasons correspond to Fall 2003, Spring 2004, Fall 2004, Spring 2005.

[18] Given the success of the show the previous host, Paolo Bonolis, successfully negotiated a multimillion contract with a different broadcaster.

share of 40 percent in prime time. This leaves no doubt as to whether the participants observed the structure of the show and were aware of the choices involved in the game.

We collected information about the participants selection by interviewing one of the three authors of the game. Before the show started the game was unknown and no ads had been run. Endemol contacted several associations and cultural groups that might suggest the name of possible participants. After the first season candidates started applying and selections were made from the pool of applicants. The interview was limited to a few minutes and did not involve any questions regarding risk attitudes or previous gambling experiences. Throughout the four seasons the only two criteria followed to select participants were telegenic look and 'average' income. According to the first criterion people were selected who showed an entertaining personality and/or looks. According to the second criterion participants that were believed to have high incomes were rejected. As one of the authors puts it: "We would not select the wife of a jeweler". While we believe the former criterion does not create severe selection bias, the latter criterion might make the sample not representative of the population. This feature of the selection also limits the extent to which we can compare risk attitudes across income levels in this study. As we admit this potential problem, we should report that the intention of the authors was to select the "average Italian" so that the audience could relate to the participant. However, based on observables, participants appear close to the median of the general population, with some evident difference only in the over-representativeness of small regions (both relatively rich, like Valle D'Aosta, and relatively poor, like Basilicata and Molise).

Self-selection is a concern in this set-up where participation is entirely voluntary and where there is certainly a cost for participating. In particular one might suggest that the show might attract relatively more risk-loving subjects and this would bias the estimate of risk aversion downwards. To address this concern, we asked the authors what seemed the main motivation behind the participants' applications. Many reported the "desire to be on TV" as the main reason for applying, followed by "the need for some cash". Moreover, as described above, participants were contacted by the producers in the first season. This allows us to test whether self-selection is an issue in the subsequent seasons employing a contrast of time indicators or considering the presence of time trends in our estimates. Self-selection would also be a concern if this show required a specific skill or a relatively high level of education. The lack of requirements along this dimension makes self-selection not an issue in this respect.

3 The Model

This section models the structure of the contestant decision problem explicitly, as required by the estimation procedure. In order to present the individual problem concisely we describe two nested alternatives concerning the informational content of the offers. Under the first alternative offers reveal nothing concerning the content of the box held by the contestant at the time of the offer. Under the second alternative we discuss the case where offers have informational content. The contestant optimally makes use of such information. Finally we present some remarks concerning the role of the offerer.

3.1 The contestant's decision problem

3.1.1 Preferences

For each game show played by a contestant $t = 1, ..., T,$[19] let $u_t(y_t)$ indicate t's vN-M utility function where y_t represents t's wealth and the contestant is assumed to behave as an expected utility maximizer. We allow for heterogeneity in the contestants' preferences for risk and parameterize $u_t(y_t)$ by a random vector θ_t joint normally distributed $N(M_t, \Sigma_t)$ and independently across t's. An (important) benchmark example we will employ throughout is $u_t(y_t)$ being constant relative risk aversion (CRRA), i.e. $u_t(y_t) = y_t^{1-\gamma_t}/(1-\gamma_t)$, with $\gamma_t \sim N(\mu_t, \sigma_t)$.

3.1.2 The optimization problem

Let us consider the problem of contestant t (indexes are omitted). Define a stage s of the game as the phase of the game when an offer is made and the contestant has to decide to accept or reject the offer. From the description in Section 2 we set $s \in \{1, 2, 3, 4, 5\}$. At stage s the contestant faces a lottery characterized by a set of prizes K_s and corresponding probabilities P_s. The set $K_s = \{k_i\}$ consists of prizes not yet discovered at stage s and P_s consists of the probabilities assigned to each of those prizes being held by the contestant:

$$P_s = \{p_i | k_i \in K_s\},$$

with $\sum_i p_i = 1$, so for instance at $s = 1$, $p_i = 1/14$. Denote also the (unknown) prize contained in the contestant's box as $k_i' \in K_s$.

Suppose now that at every stage s the participant is presented with a (positive) monetary offer[20], that we indicate with m_s. The contestant's decision consists of either accepting the offer, thereby terminating the game and renouncing the option value of playing, or rejecting the offer and proceeding to the next stage $s+1$. Accepting an offer is therefore an absorbing state. Indicate the decision by the contestant at stage s as $d_s \in \{0, 1\}$, where $d_s = 1$ if the offer is accepted and $d_s = 0$ if the offer is rejected.

At s the contestant compares the utility from the monetary offer m_s with the continuation value from rejecting the offer and therefore faces a dynamic optimization problem where the value function of the monetary offer V at s is the maximal:

$$\begin{aligned} & V_s(P_s, K_s, m_s) \\ = \; & \max_{d_s \in \{0,1\}} \{d_s u(y + m_s) + (1 - d_s) E_s [V_{s+1}(P_{s+1}, K_{s+1}, m_{s+1}) | \Omega_s, K_s, P_s]\}. \end{aligned} \tag{1}$$

In (1) $E_s[\cdot]$ indicates the expectation operator conditional on the information set at s, Ω_s, and the state of the game (K_s, P_s). The expectation E operates over all the possible paths implied by the sequence of box

[19] Given that each contestant plays only one game we let the index t indicate both the contestant and the game played.

[20] The possibility that the offer consists of an change of boxes is not relevant for the contestant's problem under the assumption $A1$ (below), as it will become clear in the following section. For this reason in this section we consider only monetary offers in the game. This simplifies the exposition and the analysis. The Appendix contains the formulation inclusive of the "change" option.

openings and over the distribution of offers conditional on the lottery faced by the contestant at the next stage $s+1$. Let $\tilde{h}_{s+1}$ be a discrete random variable that indicates the path followed between stage s and stage $s+1$ of the game. $\tilde{h}_{s+1}$ has support of dimension H_{s+1}, where H_{s+1} is the number of possible combinations of triplets of prizes out of the remaining set K_s. For instance, since K_4 is composed of five prizes, then the number of possible paths followed by opening any three boxes and discovering three prizes in K_4 is $\binom{5}{2} = 10$. Hence, $H_5 = 10$ and the realizations $h_5 \in \{1, 2, ..., 10\}$. Indicate with $q(h_{s+1})$ the probability of a specific path h_{s+1} being followed. Let $\tilde{m}_{s+1}$ be a (discrete) random variable distributed according to the probability distribution $f_{s+1}(m_{s+1}|h_{s+1})$. That is, $f_{s+1}(.)$ is the distribution of the offer at stage $s+1$, which is a function of the path followed between s and $s+1$, and captures the equilibrium strategy of the offerer. The discreteness assumption of $\tilde{m}_{s+1}$ is made solely for consistency with our empirical implementation.

The continuation value of the game at stage $s+1$ is given by:

$$\begin{aligned} & E_s\left[V_{s+1}(P_{s+1}, K_{s+1}, m_{s+1})|\Omega_s, K_s, P_s\right] \\ = & \sum_{h_{s+1}} q(h_{s+1}) \sum_{m_{s+1}} f_{s+1}(m_{s+1} \mid h_{s+1}) V_{s+1}(P_{s+1}(h_{s+1}), K_{s+1}(h_{s+1}), m_{s+1}) . \end{aligned} \tag{2}$$

where we use the notation $P_{s+1}(h_{s+1})$ and $K_{s+1}(h_{s+1})$ to emphasize that they depend on the actual path followed. The player information set Ω_s at s includes the distributions $q(.)$ and $f_{s+1}(.)$. The expression (2) relies on the assumption:

$A1$: *Offer $\tilde{m}_s$ follows a probability distribution $f_s(m_s|h_s)$ independent of the prize held $k_i' \in K_s$.*

Given $A1$ the probability of each path considered by the contestant is constant at any s:

$$q(h_s) = 1/H_s$$

and so are the probabilities assigned to each remaining prize:

$$P_s = \left\{ p_i = \frac{1}{n(K_s)} \mid k_i \in K_s \right\}, \tag{3}$$

where the function $n(K_s)$ simply indicates the number of elements of K_s. The expression (3) presents a trivial form of updating (i.e. counting the boxes remaining). Moreover, given that the opening of the boxes is performed by an uninformed contestant, there is no Monty Hall's problem[21] in the computation of P_s.

3.1.3 Solution and estimation of the dynamic model

The solution of the problem is numerical. Consider an offer at the last stage of the game, $s = 5$, when only prizes $K_5 = \{k_i, k_j\}$ are left and the probability of the participant having prize k_i in her box is $p_i = 1/2$. The maximal (1) can be rewritten as:

$$\begin{aligned} & V_5(P_5, K_5, m_5) \\ = & \max_{d_5 \in \{0,1\}} \left\{ d_5 u(y + m_5) + (1 - d_5)\left[u(y + k_i)p_i + u(y + k_j)(1 - p_i)\right] \right\} . \end{aligned}$$

[21] Please refer to the Appendix B for a detailed discussion of Bayesian updating in a Monty Hall type of problem and the differences from this game's structure under $A1$.

Therefore the final decision is:

$$d_5 = I\left[u(y+m_5) > (u(y+k_i)+u(y+k_j))/2\right]. \tag{4}$$

In (4), $I[\cdot]$ is an indicator function taking value 1 if the expression in brackets is true, 0 otherwise. For any possible pair of prizes in the set K_4 such decision can be calculated. It follows that (1) can be solved by backward recursion. For instance, when offer m_4 is made there are $\binom{5}{2} = 10$ possible ending pairs at stage 5 to be considered in the choice made at 4.

Within an episode a contestant is assumed to solve the model above. Data consist of her observed choices for each actual offer and the complete evolution of the prize distribution. Since the opening of boxes produces randomly independent draws and therefore the evolution of the prize distribution is exogenous, the probability of any sequence of choices for stages $1, ..., \bar{S} \leq 5$, where $\bar{S}$ is the ending stage, is given by:

$$\Pr\left[d_1, ..., d_{\bar{S}}\right] = \prod_{s=1}^{\bar{S}} \Pr\left[d_s | P_s, K_s, m_s\right]. \tag{5}$$

The likelihood contribution of the game played by a contestant t is therefore given by (5) with the addition of the appropriate indexing by t of decisions, offers, and lotteries. Hence, the log likelihood of the model for the T game shows recorded is given by:

$$\begin{aligned} \log L &= \sum_{t=1}^{T} \log\left(\prod_{s=1}^{\bar{S}_t} \Pr\left[d_s^t | P_s^t, K_s^t, m_s^t\right]\right) \\ &= \sum_{t=1}^{T} \sum_{s=1}^{\bar{S}_t} \log\left(\Pr\left[d_s^t | P_s^t, K_s^t, m_s^t\right]\right). \end{aligned} \tag{6}$$

3.1.4 Solution and estimation of the static (myopic) model

Finally, consider the problem of a static contestant, that is a myopic player who disregards the continuation value of (1) and at every stage decides according to a static rule depending exclusively upon her risk preferences:

$$d_s = I\left[u(y+m_s) > \frac{1}{n(K_s)} \sum_i u(y+k_i)\right] \tag{7}$$

where the summation is over the remaining prizes in K_s. Notice that ignoring the dynamic component of the game does not change the likelihood contribution (5) nor the log likelihood (6). Since at each stage of the game the contestant exclusively evaluates offers and lotteries according to her taste for risk, she also disregards any form of strategic interaction within the game. The following two subsections describe incorporating strategic interaction in the dynamic problem of a more sophisticated contestant.

3.2 The contestant's decision problem with informative offers

Let us now relax assumption $A1$ and consider:

$A1'$: *Offer* $\tilde{m}_s$ *follows a probability distribution* $f_s(m_s | h_s, k_i' = k_i)$, $k_i \in K_s$, where the likelihood of the signal "offer" is f_s. We assume that the (dynamically optimizing) contestant now takes into account that the monetary offer m_s received is correlated with the prize contained in the box held at the time of the offer. The probability distribution f_s is assumed identical for all t and in the empirical section we will approximate it with the empirical equilibrium play we actually observe in the data. In addition we assume that:

$A2$: *Every contestant updates her priors by Bayes rule.*

By Bayes rule it follows that the posterior probabilities assigned by the contestant to each remaining prize are:

$$P_s(m_s) = \left\{ p_i(m_s) = \frac{f_s(m_s|h_s, k_i = k_i')\, p_i(m_{s-1}, h_s)}{\sum_j f_s\left(m_s|h_s, k_j = k_j'\right) p_j(m_{s-1}, h_s)} \mid k_i \in K_s \right\} \tag{8}$$

where $p_i(m_{s-1})$ indicates the prior held before the offer m_s is made[22], but after a new triplet of prizes has dropped out[23]. Note that both (8) and (3) depend on the path followed h_s but that in general will differ under $A1'$ and $A2$.

Under assumptions $A1'$ and $A2$ the choice set of the contestant can be expanded in two dimensions. First, it is now possible to incorporate in the model the behavior of the contestant when the offer made is a "change" and not only a monetary offer. At s the probability π_s of receiving a change instead of an offer is assumed independent of $k_i' \in K_s$ and h_s (condition that we verify in the data at every stage[24]). In a change at stage s no signal is given, but the player can decide to swap boxes before continuing. The fact that monetary offers are uninformative under $A1$ and $p_i \in P_s$ are constant makes considering changes in the decision problem (1) irrelevant (since each box has the same probability of containing each prize). However, under $A1'$ and $A2$ each p_i is not constant and critically depends on the informativeness of the previous monetary offers (if any). If the probability assigned to holding a high prize in K_s is low, then by changing the box held with another one the contestant increases the chances of ending up holding a higher prize.

Second, the decision of which box to open between stages can become relevant. Consider having accepted a change at s and rejected an (informative) offer m_{s-1}. Is it optimal to open the old box the contestant was holding before the change or to open only boxes for which the posteriors are the same? The decision depends on the solution of a dynamic programming problem that extends (1) to the case of informative offers and can be found in Appendix.

[22] In case no previous offer m_{s-1} exists it holds that $p_i(m_{s-1}) = p_i = 1/n\,(K_{s-1})$.

[23] Note that Bayesian updating upon opening of a box and dropping out of prize k_r implies that $P_s(m_s, h_s'') = \left\{ p_i(m_s, h_s'') = \frac{p_i(m_{s-1}, h_s)}{\sum_{j \neq r} p_j(m_{s-1}, h_s)} \mid k_i \in K_s'' \right\}$ where $K_s = k_r \cup K_s''$ (K_s'' the set of prizes remaining from the original set and h_s'' is the updated history).

[24] Results available from the authors upon request.

3.3 Remarks on the offerer's decision problem

Some important remarks are in order concerning the role of the "Infamous" within the behavioral framework above.

First, by relying on the estimated equilibrium strategy of the offerer, $f_s(.)$ for $s = 1, ..., 5$, it is possible to abstract from modeling explicitly any component of the offerer's decisional problem in order to close the model. There are two reasons why we follow this approach. The first is simplicity. The objective function of the offerer is multi-dimensional and this complicates its formalization substantially (for instance, it depends on: the length of the episode, the amount disbursed relative to the show's per-episode production budget, the history of payments). Second, our approach seems appealing on the grounds that the multiplicity of equilibria of the dynamic game between the informed offerer (a long-run player) and the uninformed contestant (a short-run player who observes the entire history of the show) could be resolved only under relatively arbitrary refinements. Our solution is to rely on the data concerning which equilibrium is played.

In an online Appendix we propose a simple framework to represent the interaction between the offerer and the contestant and we show that the optimal strategy of the Infamous depends on the assumption about the behavior of the contestant. We assume that the offerer is risk neutral and the contestant is risk averse. We model the interaction as a reputation game with a series of short-run players (the contestants) observing the behavior of the long-run player (the offerer). We consider different strategies where the Infamous varies the offer according to the content of the box in the contestant's hands and we conclude that if the contestant updates her beliefs about the content of the box held k_i', the optimal strategy for the Infamous (the one that minimizes the expected payout to the contestant) is to offer a constant amount (the certainty equivalent of the lottery). If the contestant were risk neutral then the offerer's expected payout to the contestant would be equal to the expected value of the lottery, as the offerer would be indifferent between making an offer correlated with the content of the box held and making a constant offer. Since the contestant is risk averse, the least-cost strategy for the offerer is to fully insure her and signalling the content of the box through the offer is suboptimal as it stands in the way of full insurance. We show that it is optimal for the offerer to make an offer m that is positively correlated with the content of the box held by the contestant only when the contestant does not optimally update her beliefs based on the offer received, when she does not employ (8). In that case making a relatively higher offer guarantees acceptance on the part of the contestant and saves the Infamous the high prize. Conversely, a low offer is rejected and the low prize is paid. Although we realize that the example in Appendix postulates the total absence of Bayesian updating, we believe imperfect Bayesian updating can justify the observed correlation between the offer made and the prize held by the contestant.

Employing the empirical likelihood of the offer instead of a theoretical best response in the contestants' problem has clearly potential shortcomings. Particularly, it may act as a confounding factor in disentangling the relative role of rationality in the strategy played by the contestant and of risk preferences in the estimation procedure. In order to address this issue we employ the following strategy. We begin by estimating a behavioral model for the contestant that relies exclusively on risk preferences and abstracts from any strategic component: the static model assuming (7). We then add progressive layers of sophistication in the contestant's behavioral model (we estimate the uninformative and then the informative offers dynamic model) and observe the variation in the estimates. If estimates do not change substantially as the model assumed for the contestant approximates a more sophisticated behavior, then it is possible to argue that the estimates of risk preferences are not excessively sensitive to the behavioral assumption and the strategic component does not appear to be first order.

A final remark relates to the computational implementation. Empirically the offerer's equilibrium strategy appears to be independent on the parameterization of the individual preferences, that is $\tilde{m}_s$ is independent of the individual characteristics of the contestant. While the behavioral framework above does not rely upon this assumption in any fundamental way, it is worth noting that the estimation strategy turns out to lose substantial computational burden. We will refer to this assumption in what follows as $A3$. $A3$'s empirical content is validated in the reduced-form analysis available from the authors on request.

4 Estimation results

This section presents the estimation results for the model. Under the distributional assumption concerning individual preferences, estimation by Maximum Likelihood is feasible. This allows us to estimate the conditional moments of the distribution from which the individual risk aversion parameters are drawn. Under the normality assumption, the first two moments of the distribution of θ characterize the entire distribution.

4.1 CRRA parameterization

We begin by introducing a parametrization for the contestant t's utility over wealth. Let $u_t(y_t)$ be constant relative risk aversion, $u_t(y_t) = y_t^{1-\gamma_t}/(1-\gamma_t)$. Consistently with previous research[25], let us assume that the initial reference point for t's individual wealth y_t is her annual labor income. We do not observe annual labor income directly, but we are able to approximate it by matching the information concerning job and city of provenience provided by the contestant during the game show with the Italian bureau of census relevant sources. The relative risk aversion parameter follows $\gamma_t \sim N(\mu_t, \sigma_t)$ and in particular for observation t:

$$\begin{aligned} \gamma_t &= x_t'\beta + \eta_t \\ \eta_t &\sim N(0, \sigma_t^2) \\ \sigma_t^2 &= \exp(z_t'\delta), \end{aligned} \tag{9}$$

where both x_t and z_t are vectors inclusive of individual and time covariates. The individual covariates we are able to recover from the game show tapes are: gender, age, city and region of provenience, employment characteristics (employed or not, type of job), household characteristics (married, divorced, or not married, number of children, age and sex of children). We do not have access to any private information regarding the contestants. The time covariates include date and broadcasting season of the show indicators.

[25] See for instance Gertner (1993) and Cohen and Einav (2007).

At each stage s of each game t, it is possible calculate the continuation value of playing the game as a function of the RRA parameter γ, given all possible future box openings, the empirical distribution of future offers as a function of expected value of the lottery, the optimal behavior at future offers, the empirical probability of change versus monetary offers. The continuation value is a monotonic, decreasing, and convex function of γ and therefore there exists a unique level of γ at which the contestant is indifferent between the offer and continuing the game. If the contestant is more risk-averse than what the threshold implies (i.e. the RRA parameter γ_t falls above the threshold) she accepts the offer. If the contestant is less risk-averse, i.e. the RRA parameter γ_t falls below the threshold, the contestant rejects. This reasoning implies that the acceptance/rejection behavior on the show imposes (stochastic) bounds on γ_t: an upper bound $\gamma_{s,t,H}$ in case of a rejection or a lower bound $\gamma_{s,t,L}$ in case of an acceptance. However, since accepting an offer is an absorbing state, there is at most one $\gamma_{\bar{S},t,L}$, while there can be multiple $\gamma_{s,t,H}$ at different s. Under assumption $A3$ the process generating the stochastic bounds allows to partition the set of observations into three groups: T_R, individuals for which only rejections are observed; T_A individuals for which only an acceptance is observed; T_{AR}, observations for which both rejections and an acceptance are observed. Under this parameterization the log likelihood of the model (6) for the T game shows recorded can be rewritten as:

$$\begin{aligned}
\log L &= \sum_{t\in T_{AR}} \ln\left[\Phi\left(h_t-\mu_t^*\right)-\Phi\left(l_t-\mu_t^*\right)\right] \\
&\quad + \sum_{t\in T_R} \ln\left[\Phi\left(h_t-\mu_t^*\right)\right] + \sum_{t\in T_A} \ln\left[1-\Phi\left(l_t-\mu_t^*\right)\right] \\
\text{and } \gamma_{t,H} &= \min\left\{\gamma_{s,t,H} \middle|\ s>2\right\}
\end{aligned}$$

where we define $h_t = \frac{\gamma_{t,H}}{\sigma_t}$, $l_t = \frac{\gamma_{t,L}}{\sigma_t}$, $\mu_t^* = \frac{x_t'\beta}{\sigma_t}$ and where for computational feasibility we restrict our attention only to the final 3 stages of the game[26].

The computation of the stochastic thresholds differs between the uninformative and informative offers case. While the contestant employs constant probabilities in the former, she employs the equilibrium strategy of the offerer in the latter. We obtain the empirical signal (offer) likelihoods from the data, for each stage $s > 2$ and for each box held by the participant at stage s (relative to remaining boxes). We discretize the support of the offers in terms of percentages of expected value (due to data restrictions we allow for 5 segments: 0-20, 20-40, 40-60, 60-80, 80-100) and obtain the empirical frequency of offers at stage s as a function of the box held by the contestant (three likelihood matrices, one for each stage). The informativeness of the offer also varies with the expected value of the lottery and empirically is the strongest for lotteries with mean above $20,000$ euros. We therefore model the signal as uninformative below such amount. The contestant thus performs Bayesian updating of her priors conditional on the offer observed at every stage

[26] Empirically we observe always a change at $s = 1$ (with one exception, an offer of irrelevant amount that was rejected). At this stage such decision is always uninformative, as the player begins the game with flat priors (we verify empirically that actual randomization of the boxes' content is implemented). We exclude stage 2 from the analysis for two reasons. First m_2 is always a monetary offer of irrelevant amount as the producers of the game try to avoid an excessively premature ending. Second, the computational burden for this early stage is extremely intense. However we choose to employ the decisions at stage 2 in two particular instances where substantial inconsistencies in the choice set were identified.

and along all future lottery and possible offer paths. The empirical likelihoods are reported in Table A3 in the Appendix.

Finally, notice that the computation of the stochastic thresholds is the simplest in the static contestant's problem, that is an individual deciding on the basis of (7), as option values are disregarded.

4.2 CRRA results

4.2.1 Reduced-form calculations

To build intuition on the quantitative implications of our model we begin by briefly considering an extremely reduced-form calculation. Table 2A reports the expected value of the initial lottery of Table 1A. It is $52,560$ euros. Suppose all the contestants do in this game show is to buy insurance. Then, given that they take home an average final payoff of $29,436.09$ euros, what is the CRRA parameter estimates such that this amount is the exact certainty equivalent? For a representative contestant with annual labor income around $18,000$ euros (the mean of our sample[27]) it is $\gamma = 0.53$. The certainty equivalent for a $\gamma = 1$ contestant would be instead lower, around $18,000$ euros. For a contestant employing as a reference point for wealth a ten percent annuity of her future annual labor income (i.e. a lifetime wealth around $180,000$ euros) calculation shows that $\gamma = 1.62$. A mean parameter of constant relative risk aversion around 1 is an estimate that concords with the majority of the theoretical literature, see for instance the discussion in Weitzman (2007). This result is also in accord with recent empirical calibrations derived from labor supply decisions such as those reported in Chetty (2006). We will keep these back-of-the-envelope calculations in mind when discussing the structural estimates below, as we would be wary of large discrepancies from these intuitive results.

Another coarse, but intuitive way to look at the data is to focus on the fifth stage of the game. All the contestant does at this point is to choose between a monetary offer m_5 and a lottery consisting of two prizes. Abstract from any sophistication in the updating of the subjective probabilities and assume that the average contestant finds the two prizes with probability 1/2. What is the CRRA parameter that maximizes the score of correctly predicted decisions in the observed sample? Table 2B indicates that out of the 78 episodes reaching $s = 5$ a parameter $\gamma \in (1, 2)$ performs the best in matching actual decisions. Again a quantity that seems reasonable. For the interested reader Table 2C reports a random sample of the actual and predicted decisions (about 1/3 of the overall data available).

4.2.2 Structural-form static results

Table 3A presents the maximum likelihood estimates of the risk preference distribution for the static model (7). This is the simplest behavioral model for an individual decision-maker, as it relies exclusively on risk preferences. We employ three different specifications of the conditional density for unobserved taste heterogeneity. In column (1) we begin by estimating a model where risk preferences γ_t are normally i.i.d.

[27] The Purchasing Power Parity over GDP for Italy in 2000 as reported by the Penn World Tables (Mark 6.1) is $1,709.11$ Italian lira with and the lira/euro parity is $1936,27$. This implies a PPP adjusted average annual labor income of $18,000$ euros to be equivalent to approximately $20,392.40$ US dollars.

with mean μ and standard deviation σ, that is both x and z are assumed constant. The mean CRRA parameter is estimated at 0.8834 with an asymptotic standard error of 0.0473 while the standard deviation is estimated at 0.5341 with an asymptotic standard error of 0.0329. Both estimates are very precise and indicate that a log utility approximates the mean preferences of our contestants quite well. Particularly, the estimate for the mean accords to sensible theoretical predictions and to common perception in the literature. At the same time we uncover substantial variation in taste heterogeneity for the population, a dimension that has received limited attention in the empirical research of choice under risk[28].

In column (2) we allow the conditional mean of the unobserved risk parameters to depend on a set of observable characteristics of the individual, the contestant's income level, and a set of time covariates. Beginning with the set of individual covariates we are able to test if risk preferences differ with respect to a polynomial in age, gender, or region of provenience. We do not observe systematic differences in terms of age, a result common to the majority of the literature. The F-test p-value for joint significance of the age polynomial is however not very high: 0.11. More interestingly, we do not find female contestants to present levels of constant relative risk aversion different from male players. If anything, the opposite. Albeit part of the literature seems to support the view of a gender differential in risk attitudes, results are either mixed or not robust to alternative specifications (for instance in Cohen and Einav (2007)). A standard asymptotic chi-square test also cannot reject a null joint effects of regional dummies on the mean. A test on the coefficient of individual income cannot reject a null correlation between unobserved risk preferences and income. This is a reassuring result, given the constant relative risk aversion assumption under which we are operating.

Consider the coefficients for the time trend (based on the date the show was broadcast) and the coefficients on the seasonal indicators, as they have a specific interpretation in terms of the contestants' selection process[29]. Arguably, the sample of contestants we employ is more representative of the general population than samples of participants in laboratory experiments, but it is still necessary to produce a test of the hypothesis that our sample is approximately representative of the overall population. Indeed, it would not been unreasonable to argue that more risk-loving individuals could self-select on the game show. For this purpose we rely on a shift in the selection process with which contestants were admitted to the show. The show organizers indicated to us that the selection of the participants changed during the first season from a process in which they would be inviting members of various cultural associations to one in which people watching the show on TV would call and ask if they could participate. At that point in time, if significantly more risk-loving individuals were beginning to self-select into the pool of participants, we should observe a negative and significant trend in the levels of risk aversion of the contestants. This would imply a corresponding negative dynamic on the coefficients of the seasonal indicators[30]. This is not observed in the data. The coefficient on the time trend is positive (12.7240) and insignificant (s.e. 16.7935), while no clear pattern

[28] The exceptions in this respect include Cohen and Einav (2007) and Harrison, Lau and Rutström (2005). However the results in Cohen and Einav (2007) are particularly puzzling, even with respect to the mean of γ. Their estimates exceed eighty for the CRRA case.

[29] Also see Section 2 for this point.

[30] We thank Luigi Zingales for suggesting us this additional test.

is recognizable in the (jointly insignificant) seasonal dummies. Based on these tests we do not find evidence of preferences-based selection in the data.

For column (2) calculation of the predicted mean of the model confirms the results from column (1) ($\hat{\mu} = 0.8843$), as does the estimated standard deviation of the distribution (still precisely estimated at 0.4633 with a s.e. of 0.0434).

In column (3) of Table 3A the model maintains the same specification for the conditional mean as column (2) but now we allow for the conditional standard deviation of γ_t to depend on observables. Previous results for the mean of the RRA parameters are substantially confirmed. With respect to σ_t the main finding is a mild difference in the dispersion of risk preferences with respect to the trend. The trend variable has a negative coefficients, suggesting a reduction in the dispersion of preferences over time and is significant at the five percent level. The lack of robustness of this result to further specification checks prevents us from giving it a clearer interpretation. In general we conclude that, albeit substantial dispersion of preferences for risk is present (the average predicted $\hat{\sigma}_t$ is 0.4161), no single individual covariate seems to account robustly for dispersion in tastes.

4.2.3 Dynamic results and Bayesian updating of the contestant

Table 3B presents maximum likelihood structural estimates for the uninformative offer dynamic model. We follow the same outline of Table 3A. In column (1) the estimated mean of γ increases to 1.19 with asymptotic s.e. 0.10. Although the parameter is precisely estimated, we cannot reject that it is significantly different from the static model. Thus, it accords with the intuitive reduced-form calculations we performed. Furthermore, though statistically different, the estimates are not economically different, at least in light of recent discussion in the empirical literature concerned with estimating RRA. Chetty reports a mean value of 0.71, bound between 0.15 and 1.78 across samples, and at the upper bound for complementarity in his model he estimates $\gamma = .97$ –"a modest rise" from an estimate of 0.71 in his words. Attanasio, Banks and Tanner (2002) commenting on quantitatively similar coefficients write that "the point estimate of" the parameter of RRA (1.366) is very close to unity in the sample of likely shareholders defined with the fixed cutoff and a bit smaller (1.027) in the sample with the variable cutoff". An estimated standard deviation of 1.05 with s.e. 0.055 is however larger relative to the mean than in the static case.

Table 3C presents maximum likelihood estimates under assumptions $A1'$ and $A2$. Extending the model in this direction allows to account for the informational content that we empirically observe in monetary offers. Similarly to column (1) of Table 2, in column (1) we assume the RRA parameters to be i.i.d. The estimate of the mean of γ is 1.3688 with a s.e. of 0.1254 and the estimate for the standard deviation is 1.2800 (s.e. 0.0842). Again both estimates are significant, below the 1 percent critical level. In addition, $\hat{\mu}$ is larger but not statistically different in a chi-square test from the estimate of Table 3B for the same specification.

Concerning columns (2) and (3), the estimates confirm the findings of Table 3A-3B both in the case where the mean of γ_t depends on observables and in the case where both mean and standard deviation do. Again we reassuringly find neither systematic effect of income on the mean, nor evidence of selection. In fact, although the season dummies become significant in columns (2) and (3) for the mean, there is no clear path of self-selection detectable. For instance, in Table 3C, column (2), the season 1 dummy presents an estimated coefficient of 1.6183 with s.e. of 2.0604, season 2 is 1.5101 (s.e. 1.3789), season 3 -0.0346 (s.e. 0.6541). But for column (3) coefficients are respectively: 1.3805 (s.e. 0.9277), 1.6349 (s.e. 0.6245), 0.4120 (s.e. 0.3533).

Remark on testing for Bayesian updating. Even if monetary offers are informative, it is not necessarily the case that contestants employ Bayesian updating when deciding under risk. Is it possible to discriminate from the data if it is the case? We have already observed how this is unfeasible in presence of monetary offers alone. However, under the assumption of informative offers it is possible to test for Bayesian updating by observing if box-changing behavior conforms to theory. If for any constant relative risk aversion parameter contestants decide to change their box following informative offers, this would suggests that indeed their subjective probabilities have been updated incorporating such informational content.

A striking feature characterizes the data in this respect. Offers preceding changes are generally uninformative. Conditional on a change at $s' > 2$ the null $f_s(m_s|h_s, k_i' = k_i) = f_s(m_s|h_s)$ for the offer likelihoods cannot be rejected at any standard statistical levels. Our explanation is the following. Changes are random for given history of the game, but for games that lose all large prizes at the onset changes become more likely. Since offers are uninformative at relatively low stakes, this makes Bayesian updating untestable along this direction.

4.2.4 Testing informational assumptions

It is feasible to compare the dynamic decision models with respect to informational assumptions (dynamic with uninformative offers versus dynamic with informative offers). As the two models are non-nested, the typical econometric methodology in such instance is to run generalized likelihood ratio tests of model selection. Particularly, we employ Vuong (1989) and Clarke (2003) model selection tests. The null for both tests is that both models are true against a one-sided alternative that only one is true. The advantage of reporting both the Vuong and the Clarke tests is completeness. The former has better power properties when the density of the likelihood ratios is normal, while the latter is more powerful when this condition does not hold. The baseline specification is the informative offer model, column (1) of Table 3C, and it is tested against the dynamic model without informative offers estimated in column (1) of Table 3B. The p-values of both the Vuong and the Clarke sign tests do not favor the informative offer model with p-values above 0.9. Specifically, the Vuong p-value and Clarke sign test p-value are respectively 0.97 and 0.99.

Repeating the tests for the informative offer model against the static model yields similar results. The p-values of both the Vuong and the Clarke sign tests do not favor the informative offer model. The Vuong p-value and Clarke sign test p-value are respectively 0.58 and 0.96. Overall, the evidence seems to substantially downplay any sensitivity of the estimates with respect of the peculiar informational advantage of the Infame.

4.2.5 Static and dynamic results without offerer's information

This subsection further probes the sensitivity of our analysis to the informational assumptions. The success of the game protocol has prompted its popularity in several countries, among them the United States and the Netherlands. Data for the American and Dutch game shows have been collected among the others by Post et al. (2008), whose data we employ here. Tables 3D and 3E present the maximum likelihood structural estimates for both the static and the uninformative offer dynamic model for two samples of 47 episodes each from the U.S and the Dutch versions of the game. The main advantage of employing these data is the absence of any strategic role for the offerer, who has no information on the content of the contestant's box, while the game maintains a near identical format relative to the Italian version[31]. One should view the evidence in this subsection as an important robustness check of the validity of our analysis with respect to the informational structure of the game (hence the dropping of the informative offer model). We also interpret Tables 3D and 3E as an informative check of external validity outside the Italian experiment.

Table 3D reports the static results for both the U.S. and Dutch samples. Table 3E reports the dynamic model results for both samples. In both samples and models the main findings reported in the Italian data are by and large confirmed, including CRRA parameters not quantitatively distant from 1 and a reassuring performance of expected utility theory in terms of number of consistent choices detected (practically complete consistency -see discussion below). In the specification where risk preferences γ_t are normally i.i.d. with mean μ and standard deviation σ, (i.e. both x and z constant), all mean CRRA parameters are precisely estimated. The mean γ_t is estimated at 0.25 for the U.S. myopic model and 0.64 for the U.S. dynamic model. The data also indicate substantial heterogeneity, with estimates of the standard deviation $\hat{\sigma}$ around the size of the mean γ_t. The mean γ_t is estimated at 0.38 for the Dutch myopic model and 2.00 for the Dutch dynamic model (the latter number is driven by few early rejections at stage 3, for which the dynamic model is arguably inaccurate due to the large state-space dimension[32]). Parameter heterogeneity appears substantial within the Dutch sample as well, casting doubt on the validity of inference assuming homogenous preferences[33].

Finally, chi-square tests reject equality of the estimated means γ_t across models and samples, but cannot reject for both models and for both samples inclusion of the mean γ_t estimated in the $[0.15-1.78]$ interval

[31] We employ episodes from the Dutch and US games characterized by the same structure of box-opening sequences (this is the bulk of Post et al. (2008) data -US and NL1a samples- where the choice was driven by ease of implementation). The contestants begin with 26 boxes (not 20 as in the Italian version of the game) and at each of the 9 stages (instead of 5) they open 6, 5, 4, 3, 2, 1, 1, 1, and 1 box respectively. The prize set for the Dutch game show is [0.01, 0.2, 0.5, 1, 5, 10, 20, 50, 100, 500, 1000, 2500, 5000, 7500, 10000, 25000, 50000, 75000, 100000, 200000, 300000, 400000, 500000, 1000000, 2500000, 5000000] and for the US one is [0.01, 1, 5, 10, 25, 50, 75, 100, 200, 300, 400, 500, 750, 1000, 5000, 10000, 25000, 50000, 75000, 100000, 200000, 300000, 400000, 500000, 750000, 1000000]. Absent information about individual income we impute an income of $25,000$ euros for the Dutch contestants and $39,000$ dollars for the U.S. contestants, the relevant median income figures in 2005. Due to the large combinatorics at early stages, the U.S. game's individual thresholds are estimated from stage 4 onwards, for the Dutch game from stage 3 onwards (in the latter sample there are 10 episodes finishing at stage 3).

[32] Excluding Stage-3 rejections within the dynamic model delivers a mean RRA of 1.42 and dispersion of 1 in the Dutch sample. Stage-3 rejections do not occur in the US sample.

[33] Post et al. (2008).

indicated as reasonable by Chetty (2006).

4.2.6 Lifetime wealth measures

In the empirical analysis up to this section we have employed average labor income as an approximation of the individual's wealth. We propose two alternative reference points in order to further investigate the robustness of our results.

First, we employ, as a measure of wealth, average labor income multiplied by a factor of 10. This is reasonable if the individual discount rate is relatively low. For a given concave utility function, by increasing the reference point we necessarily consider a range over which the utility function is relatively less concave. Therefore, in order to explain the choices observed in the data, the utility function needs to be more concave than the one we estimate for the case of a reference point equal to annual labor income. For the specific case of the CRRA utility function this intuition is confirmed as we observe an estimate of γ around 3 for the three models presented (static, uninformative offer, informative offer).

Second, we employ zero wealth as a reference point. Contrary to the previous case, we are now considering a lower range of values over which a given function will be relatively more concave. Therefore, in order to generate the observed set of choices, the degree of concavity must be lower than the case of annual labor income as reference point. Indeed, the CRRA coefficient γ is precisely estimated at around 0.5 for all three models considered.

The main result in this section is that we obtain mean values of the CRRA parameters in a sensible, lower single-digit range under diverse assumptions relative to wealth. Notice also that the log likelihoods in columns (1) of Tables 3A-C are substantially higher than the ones reported in Table 4 for quite comparable number of observations, favouring our initial choice of reference point[34].

4.3 Consistency of the choice set

The unobserved heterogeneity assumption under which we operate requires consistency in the choice set of each individual, a condition that we now discuss. Table 5 reports the number of inconsistent thresholds found in the analysis of the static, the uninformative, and the informative offer model respectively. Each behavioral model provides upper bounds, lower bounds, or both, for each episode of the game show. A pair of inconsistent thresholds indicates an empty interval of RRA parameters, i.e. two choices by the same contestant implying an upper bound for rejection lower than a lower bound for acceptance. Out of a total 252 episodes, the number of games in which we find inconsistencies is 11 for the static model, 8 for the uninformative offer model, and 15 for the informative offer model. The number of inconsistent choices is always smaller than 6 percent of the total number of registered choices over monetary offers in the sample[35].

[34] The differences in the number of observations are due to inconsistencies in the choice sets (discussed in the following section) and the presence of 4 outliers, 1 in the static and 3 in the informative offer model.

[35] Even if compared to the number of episodes presenting both acceptance and rejections (across the various models we register 86, 90, and 86 "interval" observations that are consistent, as reported in Tables 3A-C) such numbers are not large. Notice however that the correct reference point is the total number of episodes, not the total number of "interval" episodes

We can distinguish inconsistencies further, depending on the distance between the acceptance and the rejection thresholds. Let us define a large inconsistency as an acceptance lower bound higher of more than 1 CRRA unit than the lowest rejection upper bound. Table 5 reports 4 large inconsistencies for the static model (8 inconsistent thresholds), 3 for the uninformative offer model, and 11 for the informative offer model. Hence, a substantial fraction of the inconsistencies are mild behavioral discrepancies.

In the U.S. and Dutch samples employed in Tables 3D and 3E the number of inconsistencies is also very low. Out of 47 episodes, the US static model indicates 6 observations displaying inconsistent choices. The number falls to 5 in the U.S. dynamic model. In the Dutch static model the number of inconsistencies is even lower, 3 out of 47 episodes and 4 in the uninformative offer dynamic version. Notice that we classify as inconsistent even observations with tiny differences, such as a threshold for rejection of .045 (the lower bound) and for acceptance of .042 (the upper bound)[36], which might well be considered not particularly damning for EUT. For instance, the 4 inconsistencies of the dynamic model in the Netherlands sample are on average 0.18, a minimal deviation.

There are two important implications of this result. First, the number of inconsistencies is sufficiently small to justify our modeling strategy based on unobserved heterogeneity. Random utility models could incorporate easily the inconsistent observations by allowing for some random "error" in the execution of the contestant's decision, of which it would be possible to estimate the distribution. This would increase the computational burden substantially however. By first solving for the stochastic bounds on the risk preferences' parameters and then estimating an interval regression model we avoid solving for the stochastic bounds in the dynamic problem at every iteration of the ML estimator. Under unobserved heterogeneity and assumption $A3$, MLE is consistent, asymptotically efficient and asymptotically normal. Had the threshold-finding routine embedded in the ML optimization, the CPU-time could become almost prohibitive (a single iteration for the informative offer model exceeds 12 CPU-hours on a fast microcomputer).

Second, the data do not seem to reject contestants being expected utility maximizers, quite the contrary especially in the static and in the uninformative offers model. Assuming that individuals decide based exclusively on their marginal utility of wealth seems sufficient to explain observed behavior in a simple and consistent matter. Expected utility is sufficient to provide a plausible representation of the individual choice under risk without producing paradoxical behavior.

4.4 A discussion of Rabin (2000)

A specific feature of this game show is its substantial variety in monetary payoffs, some of which are orders of magnitude apart. Although every episode begins with the same set of prizes reported in Table 1A, the stochastic evolution of the game produces wide differences in the set of lotteries each player faces later on in the show. Some players see small prizes dropping out early in the game and end up facing extremely attractive lotteries. Other contestants reach the final stages of the game facing lotteries consisting of trivial

since no restriction on the behavior of the contestant is imposed at the onset of the game.

[36] Dutch episode of 4/12/05.

amounts of money if compared to their lifetime wealth[37]. We exploit such variation to investigate the presence of substantial behavioral deviations from expected utility theory. The main result of this section is that contestants on the show seem to behave as expected utility maximizers.

Rabin (2000)'s calibration theorem states that an expected utility maximizer who rejects a small actuarially favorable gamble for a wide range of initial wealth levels, will reject almost any large bet and therefore exhibit an unrealistically high degree of risk aversion over large-stake gambles. This is because local concavity of the utility of wealth implies that marginal utility is decreasing at a paradoxically fast rate. The theorem has three related consequences, that we address. The first implication is that a common underlying concave utility of wealth cannot explain attitudes toward risk both over large gambles and small gambles. The second implication is that, if we attempt to employ expected utility theory to explain such attitude, we are bound to find higher levels of risk aversion at small stakes than at large stakes. The third implication is that, in order for an expected utility maximizer to have a reasonable degree of risk aversion over large gambles, she must be almost risk neutral over small gambles.

The first issue was addressed in the introduction where we showed that a logarithmic utility function does a reasonably good job in explaining the acceptance and rejection pattern across different gamble sizes.

In this section we address the second and the third implications of the calibration theorem and we find that over small gambles the estimates of the coefficient of risk aversion is smaller than over large gambles and, more specifically, utility is almost linear at small stakes.

Table 6A-C report results concerning our basic i.i.d. specification for unobserved heterogeneity in risk preferences in the CRRA case (i.e. the specification in column (1) of Tables 3A-C) at large and small stakes. With "large-stakes" we indicate a decision made in the presence of a lottery implicating a substantial departure from lifetime wealth. We consider to be a large-stake lottery one that includes at least one prize in excess of $250,000$ euros (indicated as max option) or, alternatively, a lottery with mean payoffs of at least $75,000$ euros (mean option). Conversely, with "small-stakes" we indicate a decision made in the presence of a lottery implicating at best a fairly modest departure from lifetime wealth. For instance a lottery including no prize in excess of $25,000$ euros (min option) or, alternatively, a lottery with mean payoffs of at most $18,000$ euros (mean option). For each specific behavioral model we analyze (static, uninformative, and informative offer) we sample exclusively those decisions made under the specified definitions. Each table presents the benchmark results of column (1) in Tables 3A-C (indicated as "All stakes"), and results in each of the four subsamples.

We report two results. First, the parameters of the distribution of γ can be estimated precisely by employing decisions made at large stakes only and a constant relative risk aversion parameter of about one captures the average of the sample population and its standard deviation precisely. Second, in all specifications restricted the small-stake lotteries it is not possible to reject the null of risk neutrality at the mean ($\mu = 0$) at any standard confidence level. Notice that this second result confirms the well-known proposition that expected-utility maximizing risk averse individuals should be locally risk neutral, which is

[37] Consider however that 50 Euros (the minimum payment the contestant receives) is in the high range of payoffs offered in laboratory experimental settings.

also at the basis of Rabin (2000) calibration theorem.

A reasonable criticism to the subsampling approach we follow is that large-stake lotteries are more likely to appear earlier in the game while small-stake lotteries are typically limited to latter stages $s = 4, 5$. This implies selection into continuation of the game, as more risk loving individuals are potentially more likely to reach the final stages of the game. In order to address this issue we show the corresponding $s = 4, 5$ subsample results below each large-stake lottery. With the exception of the static model we in fact observe mildly lower estimates of the mean. However, we can still reject risk neutrality at the 1 percent confidence level[38] for the large-stake lotteries.

Table 7 reports reduced-form results of small and large-stake lotteries similar to those of Table 2B for $s = 5$. The aim here is no only to show what γ maximizes the score of correctly predicted choices in each subsample, but the performance of risk neutrality, $\gamma = 0$, relative to such point. While risk neutrality predicts correctly 19/32 of the small-stake lotteries vis-a-vis a 23/32 at $\gamma = 1$, risk neutrality fares much worse in the large-stake case with a relative "score" of 9/25 vis-a-vis 21/25. This table offers an intuitive illustration of why we fail to reject risk neutrality at small stakes.

This set of empirical results contradicts a crucial point of departure in Rabin (2000) and Rabin and Thaler (2001), but it is not completely unexpected. Palacios-Huerta and Serrano (2006) offer several intuitive illustrations of local risk neutrality. Metrick (1995) fails to reject risk neutrality in a relatively small-stake game show (although with the caveats we report in the introduction).

Finally, let us briefly return to the small set of inconsistent choices described in the previous section. We are interested in investigating whether such inconsistent behavior conforms to Rabin's predictions. We define a "Rabin's observation" an individual characterized by the following decisions:

i. Rejection of an offer when facing a large-stake lottery (determining a large-stake upper bound $\gamma^L_{s,t,H}$);

ii. Acceptance of an offer when facing a small-stake lottery (determining a small-stake lower bound $\gamma^S_{s,t,L}$);

iii. The large-stake upper bound is lower than the small-stake lower bound $\gamma^L_{s,t,H} < \gamma^S_{s,t,L}$.

Among the 8 inconsistencies found in the previous section for the uninformative offer case, we find 3 "Rabin's observations"[39]. Therefore the anomaly raised by Rabin seems to be relevant in 3 out of 43 (35 consistent small-stake, min option lotteries + 8 inconsistent ones) instances where individuals face first large-stake lotteries and subsequently small-stake lotteries, all situations where such anomaly could potentially arise.

5 Alternative theories

In this section we explore a few theories that have been proposed as alternative to expected utility theory and attempt to evaluate their ability to explain the behavior exhibited by contestants on this game show.

[38] Notice also that sample sizes between the large and small-stakes lotteries become more comparable when restricting the former to $s = 4, 5$.

[39] An example of an observation of this type is given in Table 2C for the episode dated 3/12/2004.

We present formal econometric (nested and non-nested) tests that allow to discriminate across the different models.

The first candidate is Rank-Dependent Expected Utility (RDEU) introduced by Quiggin (1982), which introduces weights in the decision process that differ from the probabilities of the individual outcomes.[40] While we keep a CRRA specification for the value function[41], we consider several weighting functions $\varphi(.)$ mapping the cumulative probabilities into the interval $[0, 1]$. The first is a simple power function (RDEU1):

$$\varphi(p) = p^{\alpha}$$

The second is the function proposed by Prelec (1998) (RDEU2):

$$\varphi(p) = \exp\left[-(-\ln p)^{\alpha}\right]$$

and the third was introduced by Camerer and Ho (1994) to account for a change in concavity that would allow for the overweighting of small probabilities (of large gains) and underweighting of large probabilities (of small gains) (RDEU3):

$$\varphi(p) = \frac{p^{\alpha}}{\left[p^{\alpha} + (1-p)^{\alpha}\right]^{\frac{1}{\alpha}}}$$

It is worth emphasizing that, although we test for non-linearity of the weighting function, our framework is one in which anomalies such as the ones found by Allais (1953). This is because, first, all outcomes are equiprobable at a given stage in the game and, second, whether probabilities are large or small is not correlated with the size of the prize, but only depends on which stage of the game one is considering (2, 5 or 8 boxes left).

The next alternative we consider is Cumulative Prospect Theory (CPT) introduced by Tversky and Kahneman (1992) as a modification of Kahneman and Tversky (1979). Differently from expected utility theory, CPT postulates that individual behavior under risk depends on the evaluation of gains and losses with respect to a reference point, and not of the final wealth levels. CPT also introduces different value functions for gains (a concave value function) and losses (a convex value function) and posits that individuals also exhibit loss aversion, that is the value function is steeper in the loss domain than the gain domain[42]. As discussed above, prospect theory and later CPT and RDEU introduce weighting functions through which an individual transforms the cumulative probabilities of a lottery into weights[43]. Although in this framework there are literally no losses and the contestant can only win positive sums, we set our reference point to be

[40] Although Kahnemann and Tversky (1979) introduced weighting functions, Quiggin (1982) proposed applying such functions to cumulative probabilities to restore first-order stochastic dominance, which was violated by prospect theory.

[41] The value function we consider is $v(k) = k^{1-\gamma}/(1-\gamma)$ for $\gamma > 0$ and $\gamma \neq 1$, and $v(k) = \ln k$ for $\gamma = 1$.

[42] Blavatskyy and Pogrebna (2006b) propose to employ change offers to deter violations of loss aversion. They interpret the acceptance of a change offer as a violation of loss aversion since all outcomes are equiprobable. At the first stage 251 change offers were made and they were rejected in 196 cases. At the second stage the frequency of rejections was 1/5; at the third stage 14/24; at the fourth 18/25; at the fifth 46/80.

[43] Editing rules for the decision weights and memory of the decisional process are also typical dimensions within the prospect theory framework from which we abstract. See Thaler (1991) for a discussion. Our approach is similar in this respect to the one followed by Jullien and Salanié (2000).

offer at any given stage. We then calculate the gains and losses as a difference between each possible prize and the offer. This seems like a reasonable reference point since the offer is a sure amount of money the contestant could walk away with and any amount less than that can be considered a loss of that certain sum.

In order to reduce the computational burden of estimating many parameters, we adopt linear value functions for both gains and losses, but allow for loss aversion. The value function for gains is:

$$v_G(x) = x$$

while the value function for losses is:

$$v_L(x) = \lambda x$$

As for the weighting function we consider the three alternatives described above: power function (CPT1), Prelec (1998) (CPT2) and Camerer and Ho (1998) (CPT3).

Again, in order to reduce computational burden, we consider the static version of the game but extend the analysis to stages 2 to 5 of the game (stage 1 being generally uninformative). Specifically, consider a static decision maker similar to the one illustrated by (7) under expected utility. The stage-decision under RDEU and CPT can be represented as:

$$\text{RDEU} \quad : \quad d_s = I\left[v(m_s) > \sum_i \varphi(p_i)\, v(k_i)\right] \tag{10}$$

$$\text{CPT} \quad : \quad d_s = I\left[0 > \sum_i I(k_i - m_s > 0)\,\varphi(p_i)\, v_G(k_i - m_s) + \left[1 - I(k_i - m_s > 0)\right] v_L(k_i - m_s)\right] \tag{11}$$

where the summation is over the remaining prizes in K_s.

The procedure we follow is similar to the one employed in the expected utility case where we allow for individual-specific heterogeneity of preferences. We identify for each player an indifference function on the (γ, α) plane (for RDEU) and on the (λ, α) plane (for CPT). These indifference functions are such that for points above these functions, there is acceptance and below it the offer is rejected. We maintain the assumption of unobserved heterogeneity as in (9), a condition that we now extend to:

$$\begin{aligned} \gamma_t &= x_t'\beta + \eta_t \\ \alpha_t &= x_t'\delta + \varepsilon_t \\ \begin{pmatrix} \eta_t \\ \varepsilon_t \end{pmatrix} &\sim N\left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_\eta^2 & \sigma_{\eta,\varepsilon} \\ \sigma_{\eta,\varepsilon} & \sigma_\varepsilon^2 \end{pmatrix}\right). \end{aligned} \tag{12}$$

The stochastic formulation for parameters (λ, α) is analogous.

Table 8 reports the MLE for the static model incorporating decisions made according to (10) by individuals with preferences given by (12) for an i.i.d. specification in which the conditional means of γ and α are constant[44]. We report results for the three RDEU models (power, Prelec 1998, and Camerer and Ho 1994)

[44] We run but do not report a specification with conditional means depending on a set of covariates similar to those employed in column (2) of Tables 3A-C (with the exclusion of regional and season of the show fixed effects) because it produces very similar results.

in the three sections of the table. The estimates of the mean of γ in the i.i.d. specification is 0.68 (s.e. 0.22) for RDEU1, while the standard deviation of γ is 0.25 (s.e. 0.33). These estimates do not differ statistically from column (2) of Table 4 nor of Table 3A, given the relatively larger standard errors (all results for the mean remain significant however). Assuming RDEU2 and RDEU3 does not change radically the picture. We now obtain, however, smaller results for the mean of γ. The additional parameters of interest are the mean of α, 0.72 (s.e. 0.26), and its standard deviation estimated at 0.31 (s.e. 0.38) for power-weighting. The estimate of the covariance between the weighting function exponent α and γ is -0.07 (s.e. 0.20). For Prelec weighting the mean of α is 0.81 and for Camerer-Ho weighting is 0.66. The null that the mean of α is equal to 1 cannot be rejected at standard confidence levels tends to lend support to the claim that expected utility theory (which implies linear weights) does not sensibly underperform CPT in this setting. The null that the standard deviation of α is equal to 0 cannot be rejected at standard confidence levels tends to support the claim that no excessive dispersion in the α is disregarded under EUT.

In Table 9 we report the results for loss aversion. We find prima facie evidence of a certain degree of loss aversion. The power weighting, Prelec, and the Camerer-Ho (CPT1, CPT2 and CPT3) models present statistically significant levels of loss aversion with a mean of λ between 1.92 and 3.08. These figures are similar to the estimates discussed in Tversky and Kahneman (1992) and those in much related literature. Interestingly the estimates of the variance of the loss aversion parameter are also statistically significant at least at 10 percent confidence, indicating a certain degree of heterogeneity in preferences against losses (about 50% of the mean of λ).

Concerning weighting, Table 9 supports the conclusion of Table 8. The null that the mean of α is equal to 1 cannot be rejected at standard confidence levels for all CPT models.

Differently from Table 8, the introduction of loss aversion makes the EUT and CPT choice models non-nested. We follow Harrison and Rutstrom (2006) and Harrison (2006) in implementing Vuong (1989) and Clarke (2003) model selection tests. Once again, the null for both tests is that both models are true against a one-sided alternative that only one is true. The baseline specification for EUT is the first we presented, column (1) of Table 3A, and it is tested against each of the three CPT models estimated in Table 9. The p-values of both the Vuong and the Clarke sign tests favor EUT over each of the three CPT models with p-values below 0.02. Specifically, the Vuong p-value and Clarke sign test p-value are respectively 0.012 and < 0.001 for EUT over CPT1; and < 0.001 for both tests of EUT over CPT2 and CPT3.

6 Concluding remarks

We employ a natural experiment to estimate moments of the unobserved distribution of constant relative risk aversion parameters in the population and find precisely estimated average RRAs slightly above one. We also detect substantial dispersion in risk preferences, as we estimate the standard deviation of the RRA parameter to be around one. The introduction of parameter heterogeneity appears to be both important for capturing a realistic feature of the data and in furthering our understanding of the appropriateness of EUT. Within our estimation framework EUT performs extremely well in terms of number choice rationalized.

We also address the concern, raised by Rabin (2000), that expected utility might not be suitable to explain the different behavior of individuals facing small and large-stakes gambles. As Palacios-Huerta and Serrano (2006) point out, Rabin's calibration theorem and therefore the prediction of unrealistic levels of risk aversion over large-stake lotteries, relies on the rejection of small gambles for a wide range of wealth levels. Similarly to the results reported by Palacios-Huerta and Serrano (2006) we find that in a real-life experiment small gambles are accepted, indicating that a local estimate of the degree of risk aversion for small-stakes lotteries reveals almost risk neutrality. Since Rabin himself suggested that "almost risk-neutrality" over small gambles is the requirement for finding reasonable risk aversion levels at high stakes we do not consider our findings in contrast with Rabin's theorem. In fact, within the same experiment, we are not able to reject risk neutrality at low stakes, yet our estimates start picking up some curvature in the utility of income only when we move to larger lotteries. We believe this is an encouraging result as it suggests that the risk attitudes towards small-stake and large-stake lotteries can be derived from the same underlying utility of income.

References

[1] Andersen, Stefen, Glenn W. Harrison, Morten Igel Lau, Elisabet E. Rutstrom (2008), "Risk Aversion in Game Shows," in J.C. Cox and G.W. Harrison (eds.), Risk Aversion in Experiments Bingley, UK: Emerald, *Research in Experimental Economics*, Volume 12.

[2] Andersen, Stefen, Glenn W. Harrison, Morten Igel Lau, Elisabet E. Rutstrom (2006b), "Dynamic Choice Behavior in a Natural Experiment," Working Paper, Department of Economics, University of Central Florida.

[3] Attanasio Orazio, James Banks and Sarah Tanner, 2002. "Asset Holding and Consumption Volatility," *Journal of Political Economy*, University of Chicago Press, vol. 110(4), pages 771-792.

[4] Beetsma, Roel M. W. J., and Peter C. Schotman (2001), "Measuring Risk Attitudes in a Natural Experiment: Data from the Television Game Show Lingo," *The Economic Journal*, 111, 821-48.

[5] Berk, J. B., E. Hughson, and K. Vandezande (1996) "The Price Is Right, but Are the Bids? An Investigation of Rational Decision Theory," *American Economic Review,* 86(4), 954–70.

[6] Blavatskyy, Pavlo, and Ganna Pogrebna (2010) "Endowment effects? Even With Half a Million on the Table," *Theory and Decision* 68:1-2, 173-192.

[7] Chetty, Raj (2006) "A New Method for Estimating Risk Aversion," *American Economic Review,* 96(5): 1821–1834

[8] Cicchetti, Charles J., and Jeffrey A. Dubin (1994) "A Microeconometric Analysis of Risk Aversion and the Decision to Self-Insure," *Journal of Political Economy*, 102(1), 169-86.

[9] Clarke, Kevin A. (2003) "Nonparametric Model Discrimination in International Relations," *Journal of Conflict Resolution*, 47(1), 72-93.

[10] Cohen, Alma, and Liran Einav (2007) "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97(3), June 2007, 745-788

[11] Cox, James C., and Vjollca Sadiraj (2006) "Small- and Large-Stakes Risk Aversion: Implications of Concavity Calibration for Decision Theory," Forthcoming *Games & Economic Behavior.*

[12] Cutler David M., Amy Finkelstein and Kathleen McGarry, (2008) "Preference Heterogeneity and Insurance Markets: Explaining a Puzzle of Insurance," *American Economic Review*, vol. 98(2), pages 157-62.

[13] De Roos, Nicolas, and Yianis Sarafidis (2009) "Decision Making Under Risk in Deal or No Deal," *Journal of Applied Econometrics.*

[14] Deck, Cary, Jungmin Lee, and Javier Reyes (2006) "Risk Attitudes in Large Stake Gambles: Evidence from a Game Show," Forthcoming *Applied Economics.*

[15] Gertner, Robert (1993) "Game Shows and Economic Behavior: Risk-Taking on 'Card Sharks'," *Quarterly Journal of Economics,* 108(2), 507-21.

[16] Harrison, Glenn W. (2006) "Maximum Likelihood Estimation of Utility Functions Using Stata," Working Paper, Department of Economics, University of Central Florida.

[17] Harrison, Glenn W., Morten I. Lau, and E. Elisabet Rutström (2007) "Estimating Risk Attitudes in Denmark: A Field Experiment,", *Scandinavian Journal of Economics*, 109(2), June, 341-368

[18] Harrison, Glenn W., and John A. List (2004) "Field Experiments," *Journal of Economic Literature* 42, 1009–55.

[19] Harrison, Glenn W., and E. Elisabet Rutström (2007) "Expected Utility Theory and Prospect Theory: One Wedding and A Decent Funeral," *Experimental Economics*, forthcoming.

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

[21] Jullien, Bruno, and Bernard Salanié (2000) "Estimating Preferences under Risk: The Case of Racetrack Bettors," *Journal of Political Economy*, 108(3), 503-30.

[22] Kachelmeier, Steven J., and Mohamed Shehata (1992) "Examining Risk Preferences under High Monetary Incentives: Experimental Evidence from the People's Republic of China," *American Economic Review*, 82(5), 1120-1141.

[23] Kahneman, Daniel, and Tversky, Amos (1979) "Prospect Theory: An Analysis of Decision under Risk," *Econometrica,* 47, 263-91.

[24] Mazzocco, Maurizio (2004) "Savings, Risk Sharing and Preferences for Risk," *American Economic Review*, Vol. 94, No. 4: pp. 1169-1182.

[25] Mazzocco, Maurizio and Shiv Saini (2008) "Testing Efficient Risk Sharing with Heterogeneous Risk Preferences," Manuscript, UCLA.

[26] Metrick, Andrew (1995) "A Natural Experiment in 'Jeopardy!'," *American Economic Review*, 85(1), 240-53.

[27] Mulino, Daniel; Richard Scheelings, Robert Brooks, and Robert Faff (2009) "Does Risk-Aversion Vary with Decision-frame? An Empirical Test Using Recent Game Show Data," *Review of Behavioral Finance* 1:1-2, 44-61.

[28] Palacios-Huerta, Ignacio and Roberto Serrano (2006) "Rejecting Small Gambles Under Expected Utility," *Economics Letters* 91(2), May 2006, pp. 250-259.

[29] Post, Thierry, Martijnvan den Assem, Guido Baltussen, and Richard Thaler (2008), "Deal or No Deal? Decision Making under Risk in a Large-Payoff Game Show," *American Economic Review* 98:1 38-71.

[30] Rabin, Matthew (2000) "Risk Aversion and Expected Utility Theory: A Calibration Theorem," *Econometrica*, 68(5), 1281-92.

[31] Rabin, Matthew, and Richard H. Thaler (2001) "Anomalies: Risk Aversion," *Journal of Economic Perspectives*, 15(1), 219-32.

[32] Schulhofer-Wohl, Sam (2008). "Heterogeneous Risk Preferences and the Welfare Cost of Business Cycles," *Review of Economic Dynamics*, vol. 11(4), pages 761-780, October.

[33] Thaler, Richard H. (1991) *Quasi Rational Economics* (Russell Sage Foundation, New York.)

[34] Tversky, Amos, and Daniel Kahneman (1992) "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297-323.

[35] Vuong, Quang H. (1989) "Likelihood Ratio Tests for Model Selection and Non-Nested Hypotheses," *Econometrica*, 2, 307-333

[36] Weitzman, Martin L. (2007) "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, Vol. 97, No. 4 (Sep., 2007), pp. 1102-1130.

7 Appendix: Value functions

Uninformative offers

Consider first how the value function for the case with uninformative monetary offers (1) needs to be modified when we incorporate box changes in the contestant's problem. By introducing an exogenous probability π_{s+1} of a change being offered at $s+1$ we can rewrite (1) as:

$$
\begin{aligned}
& V_s(P_s, K_s, m_s) \\
= & \max_{d_s \in \{0,1\}} \left\{ \begin{array}{c} d_s u(y+m_s) + (1-d_s)\pi_{s+1} E_s \left[C_{s+1}(P_{s+1}, K_{s+1}) \right] \\ +(1-d_s)(1-\pi_{s+1}) E_s \left[V_{s+1}(P_{s+1}, K_{s+1}, m_{s+1}) \right] \end{array} \right\},
\end{aligned}
$$

where expectations are taken conditionally to Ω_s, K_s, P_s. Notice that since under $A1$ the contestant's priors are always flat, a "change" amounts to a stage where no decision is possible[45]. Hence, the value function for a change at s, indicated by $C_s(P_s, K_s)$, can be written as:

$$
C_s(P_s, K_s) = \pi_{s+1} E_s \left[C_{s+1}(P_{s+1}, K_{s+1}) \right] + (1-\pi_{s+1}) E_s \left[V_{s+1}(P_{s+1}, K_{s+1}, m_{s+1}) \right].
$$

Informative offers

We now derive the value functions for the problem with informative offers, starting from the value of a monetary offer at s. The value function of a monetary offer m_s is the maximal:

$$
\begin{aligned}
& V_s(P_s(m_s), K_s, m_s) \\
= & \max_{d_s \in \{0,1\}} \left\{ \begin{array}{c} d_s u(y+m_s) + (1-d_s)\pi_{s+1} E_s \left[C_{s+1}(P_{s+1}(m_s), K_{s+1}) \right] \\ +(1-d_s)(1-\pi_{s+1}) E_s \left[V_{s+1}(P_{s+1}(m_{s+1}), K_{s+1}, m_{s+1}) \right] \end{array} \right\},
\end{aligned} \tag{13}
$$

where expectations are now taken conditional to $\Omega_s, K_s, P_s(m_s), Z_s$. Note the introduction of the state variable Z_s, which we now define. The continuation values and path choices are slightly less straightforward in the case of informative offers. To see why, just note that the decision whether to open a box previously in the contestant's hands or another box that was not in her hands has to be taken into account when computing the set of possible continuation paths. The probability of ending on a specific path becomes a choice variable. The variable Z_s is defined as:

$$
Z_s = \left\{ \begin{array}{c} 1 \text{ if } (d_s = 1 \wedge m_s = change \wedge d_{s-1} = 0 \wedge \exists l < s : m_{s-l} \; informative) \vee (Z_{s-1} = 1) \\ 0 \text{ otherwise} \end{array} \right.
$$

and indicates the presence of (the first) change of boxes performed after an informative offer had been previously made. Z indicates if in the set of boxes not held by the contestant there is a box she previously

[45] Changes may happen even in the presence of unfinformative offers for idiosyncratic reasons, such as personal taste or superstition. They are however features that would require a substantial and not straightforward extension of the behavioral model. We disregard them in the present setting

hold in her hands *after* having received at least one informative offer. In a setting where $s \leq 5$ and under the assumption that monetary offers before stage 3 are always uninformative, $Z_4 = 1$ if there is a monetary offer at 3 and a change of boxes at 4[46].

At s the value function of a change is the maximal:

$$\begin{aligned} & C_s(P_s\left(m_{s-1}\right), K_s) \\ = & \max_{d_s \in \{0,1\}} \left\{ \begin{array}{c} d_s \pi_{s+1} E_s\left[C_{s+1}(\bar{P}_{s+1}\left(m_{s-1}\right), K_{s+1})\right] \\ +d_s(1-\pi_{s+1}) E_s\left[V_{s+1}(\bar{P}_{s+1}\left(m_{s+1}\right), K_{s+1}, m_{s+1})\right] \\ +(1-d_s)\pi_{s+1} E_s\left[C_{s+1}(P_{s+1}\left(m_{s-1}\right), K_{s+1})\right] \\ +(1-d_s)(1-\pi_{s+1}) E_s\left[V_{s+1}(P_{s+1}\left(m_{s+1}\right), K_{s+1}, m_{s+1})\right] \end{array} \right\}, \end{aligned} \tag{14}$$

where we define the new set of probabilities for the box at hand after a change decision:

$$\bar{P}_s\left(m_s\right) = \left\{ \bar{p}_i\left(m_s\right) = \frac{1-p_i\left(m_s\right)}{n(K_s)-1} \right\}$$

and where expectations are taken conditionally to $\Omega_s, K_s, P_s\left(m_{s-1}\right), Z_s$.

The continuation value in case of an offer at stage $s+1$ in (13) and (14) is given by:

$$\begin{aligned} & E_s\left[V_{s+1}(P_{s+1}(m_{s+1}), K_{s+1}, m_{s+1}) | \Omega_s, K_s, P_s(m_s), Z_s\right] \\ = & \sum_{h_{s+1}} q\left(h_{s+1}, P_s\left(m_s\right), Z_s\right) * \\ & \left(\sum_{m_{s+1}} f_{s+1}\left(m_{s+1} \mid h_{s+1}, k_i' = k_i\right) V_{s+1}\left(P_{s+1}\left(h_{s+1,} m_{s+1}\right), K_{s+1}\left(h_{s+1}\right), m_{s+1}\right) \right). \end{aligned} \tag{15}$$

with $k_i \in K_{s+1}$.If $Z_s = 0$ in (15), the probability of each path considered by the contestant at any s under $A1'$ and $A2$ is given by:

$$\begin{aligned} & q\left(h_{s+1}, P_s\left(m_s\right), 0\right) \\ = & \frac{\sum_{j | k_i \in K_{s+1}} p_j\left(m_s\right)}{\binom{n(K_s)-1}{3}}. \end{aligned}$$

The probability of each path is simply defined by Bayes rule whenever no "informed" change has been made before (i.e. $Z_s = 0$). If $Z_s = 1$ the contestant opens the sequence of boxes $\{x, y, z\}$ so that for all paths h_{s+1} the sequence produces probabilities of following each path such that:

$$\begin{aligned} & q\left(h_{s+1}, P_s\left(m_s\right), 1\right) \\ = & \max_{q'} \{ \sum_{h_{s+1}} q'_{\{x,y,z\}} * \\ & \left(\sum_{m_{s+1}} f_{s+1}\left(m_{s+1} \mid h_{s+1}, k_i' = k_i\right) V_{s+1}\left(P_{s+1}\left(h_{s+1,} m_{s+1}\right), K_{s+1}\left(h_{s+1}\right), m_{s+1}\right) \right) \}. \end{aligned}$$

[46] We derive this restriction for the contestant's problem in presence of informative offers directly from the data. Therefore, it is not necessary to generalize to the case where more than one change is performed after an informative offer is made.

8 Appendix: Variables description

Age = age of the participant (in years). The age that the individual declares during the game show. If no indication of age is directly given, the individual is assigned to an age bracket. Approximating from the appearance of the player we assign him or her to a more precise bracket (20-25, 25-30, 30–35, ...) or to a wider bracket (20-30, 25-35, 30-40, ...) solely on the basis of appearance. The corresponding age is the mid-point of the range. This procedure was necessary in the majority of cases (209 out of 254). This approximation procedure was required by the fact that we were not allowed to access the personal files of the players (some of the following covariates present similar instances at times).

Employed = discrete variable taking value 1 if the participant is currently employed, value -1 if the participant is unemployed, value 0 if the participant is out of the labor force (retired, student...). The job and employment status the individual declares during the game show is recorded for 217 participants. For each participant for which a job is recorded we define four dummy variables: Agriculture, Manufacturing, Construction, and Services, that take value equal to 1 if the participant is employed in any of the four sectors, 0 otherwise.

Female = dummy variable taking value 1 if the participant is female, 0 otherwise.

Income = income of the participant (in Euros). In order to assign an income level to a participant we employ two sources of information: the type of job he or she declares on the show (available for 217 out of 254 episodes) and the city/province/region of residence (region is always available and the city/province is available for 224 shows). We match the individual to the national annual labor income by type of job and then adjust for the province value added relative to the national average to proxy for local average level of income. Source of both data sets is the Italian national bureau of statistics (ISTAT). For annual labor income we employ "Tavola 4.12 ISTAT (Lavoro e Reddito)" which includes national labor income by type of job (measured in Italian Lire and averaged between white and blue collar within type of job) in 2001. The data are then converted in Euros using the exchange rate for the EU parity. We use value added data broken down by province in 2001 (measured in Euros) from Table "Tavola ISTAT - CONTI PROVINCIALI 2004 Valore Aggiunto per Unita' di Lavoro". We correct for inflation from 2001 to 2004 (2001-2 = 2.5%; 2002-3 = 2.7%; 2003-4 = 2.2%) employing "Tabella 6. Indice generale nazionale dei prezzi al consumo per l'intera collettività, NIC". For individuals with only regional provenience available we employ the average annual income in the region.

Marital status = single, married or divorced. The status that the individual declares during the game show. It is available for 217 participants.

Not-main-city = dummy variable taking value 1 if the participant resides in the main city of his or her province, 0 otherwise. The city is the individual declares during the game show is recorded for 227 participants.

Number of children = number of children of the participant as declared during the game show. Sex and age of children are also sometimes available although for a smaller portion of the sample. Number of children is available for 203 participants.

Region = set of 20 dummy variables for the region of which the participant is resident.. There are some rare cases in which the player represents the region of birth, but currently resides in another region. The list of regions in the North is: Valle D'Aosta, Piemonte, Liguria, Lombardia, Veneto, Friuli-Venezia Giulia, Trentino-Alto Adige, Emilia-Romagna; in the Center: Toscana, Marche, Lazio, Umbria, Abruzzo, Molise; in the South: Campania, Puglia, Basilicata, Calabria, Sicilia, Sardegna.

Season = set of 6 dummy variables for the show season during which the show was aired. The first season was broadcast during the fall of 2003.

Trend = time trend in number of days. 0 corresponds to 01 January 1960. Date employed for computing the trend is the date the show was aired on television, not the day it was actually recorded. The day of the recording was not available in the tapes.

TABLE 1A

Distribution of prizes in the game show "Affari Tuoi"

0.01	5,000
0.20	10,000
0.50	15,000
1	20,000
5	25,000
10	50,000
50	75,000
100	100,000
250	250,000
500	500,000*

Note: *The largest prize was different for the evening series of the game show. In each evening series episode three rounds (i.e. complete games) were played. The maximum prize is 500,000 euros in the first round, 750,000 in the second round and 1,000,000 euros in the third round. Each round is equivalent to a prime time episode. There are 3 evening shows.

TABLE 1B

Timing of the game show "Affari Tuoi"

	Host	Contestant
Start	Random box assigned to contestant from set of 20	
		Open 6 of 20 boxes
Stage 1	Offer 1 (always a change)	Yes/No
		Open 3 of 14 boxes
Stage 2	Offer 2 (always monetary offer)	Yes/No
		Open 3 of 11 boxes
Stage 3	Offer 3 (change/monetary offer)	Yes/No
		Open 3 of 8 boxes
Stage 4	Offer 3 (change/monetary offer)	Yes/No
		Open 3 of 5 boxes
Stage 5	Offer 3 (change/monetary offer)	Yes/No
		Open box in hand of 2 boxes

Note: Each offer consists of a change possibility (i.e. swapping the currently held box with one of the unopened boxes) *or* a monetary offer (i.e. a sum of money offered to abandon the game). Prizes discovered drop out from the game and cannot be won.

TABLE 2A
Back-of-the-envelope Calculation - 252 episodes, CRRA

Expected value of the lottery	52560.00
Average Final Payoff (AFP)	29436.09
Average income of contestant	17875.82
γ for which AFP is CE	0.53
CE for γ=1	18003.50
γ for which AFP is CE [income*10]	1.62

Note: Monetary amounts expressed in euros. Consider the lottery at the beginning of the game with equi-probable payoffs of [50 50 50 50 50 50 50 100 250 500 5000 10000 15000 20000 25000 50000 75000 100000 250000 500000]. All prizes nominally below 50 Euros in Table 1A are paid 50 euros.

TABLE 2B
Scores for episodes reaching stage 5 with moneatry offers at different CRRA parameters

	Number of correctly predicted decisions	Fraction of correctly predicted decisions
γ=0	35	44.87%
γ=1	54	69.23%
γ=1.5	57	73.08%
γ=1.8	59	75.64%
γ=2	58	74.36%
γ=3	57	73.08%
γ=4	52	66.67%
γ=5	51	65.38%
*Total**	*78*	*100.00%*

Note: * 91 shows reach 5 stages with monetary offers but in 13 cases both prizes have equal value (50 euros). This leaves 78 episodes.

TABLE 2C

Randomly selected episodes reaching stage 5 with moneatry offers.

Date of episode (D/M/Y)	Prize Not in Contestant's Box	Prize in Contestant's Box	Monetary offer at s=5	Contestant accepts. (1 Yes/ 0 No)	Would γ=0 accept?	Would γ=1 accept?	Would γ=2 accept?
13/10/2003	25000	10000	12000	0	0	0	0
16/10/2003	15000	50	50	0	0	0	0
17/10/2003	50	10000	150	0	0	0	0
20/10/2003	250	50	100	0	0	0	0
31/10/2003	50000	100000	75000	0	1	1	1
5/11/2003	250000	75000	100000	0	0	0	0
14/11/2003	75000	50	21000	1	0	0	1
24/11/2003	25000	50000	37500	0	1	1	1
5/12/2003	10000	15000	12500	1	1	1	1
20/4/2004	100	250	175	0	1	1	1
5/5/2004	5000	50000	25000	1	0	1	1
25/5/2004	100	50	50	0	0	0	0
29/9/2004	250	50	125	0	0	0	0
6/10/2004	750000	50000	100000	0	0	0	0
20/10/2004	250	100	200	1	1	1	1
29/10/2004	50	50000	25000	1	0	1	1
8/11/2004	250000	75000	100000	0	0	0	0
15/11/2004	10000	5000	8000	0	1	1	1
17/11/2004	100000	50	40000	1	0	1	1
3/12/2004	50	5000	1500	1	0	0	0
7/12/2004	50	15000	7000	1	0	1	1
23/12/2004	15000	20000	17500	0	1	1	1
27/12/2004	50	250	125	0	0	0	0
12/4/2005	500000	75000	150000	1	0	0	1
2/5/2005	75000	10000	32500	1	0	0	1
25/5/2005	50	250000	70000	1	0	1	1

Note: Monetary amounts expressed in Euros.

TABLE 3A

MLE: Unobserved Risk Preferences in the Static Model

CRRA	Dep. Variable: γ (1)		Dep. Variable: γ (2)		Dep. Variable: γ (3)	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN						
μ	0.8834	0.0473	0.8843	--	0.8883	--
Individual Covariates:						
Age			-0.0165	0.2431	-0.0681	0.2116
$(Age)^2$			-0.0007	0.0057	0.0010	0.0050
$(Age)^3$			0.00001	0.00004	-3.29e-6	-3.87e-5
Female			-0.0250	0.1300	0.0062	0.0959
Constant			-20.4564	27.8139	-15.6332	18.9354
Regional F.E. (p-value)			--	[0.7566]	--	[0.2904]
Income(/10,000)			0.4187	0.4037	0.2756	0.3487
Time Covariates:						
Time trend			12.7240	16.7935	10.2451	12.0591
Season of the show F.E. (p-value)			--	[0.1975]	--	[0.0688]
STANDARD DEVIATION						
σ	0.5341	0.0329	0.4633	0.0434	0.4161	--
Individual Covariates						
Age					0.4883	0.6889
$(Age)^2$					-0.0110	0.0156
$(Age)^3$					0.0001	0.0001
Female					-0.0044	0.2453
Constant					13.1877	17.8708
Income(/10,000)					0.6414	0.9319
Time Covariates						
Time trend					-13.5666	6.5326
Log Likelihood	-177.7120		-158.9505		-152.1448	
No. Thresholds*	518		518		518	
No. Obs.	237		237		237	
No. Left-censored Obs.	130		130		130	
No. Right-censored Obs.	21		21		21	
No. Interval Obs.	86		86		86	
chi-squared [df, p-value]	--		37.5200	[28, 0.1078]	39.8800	[28, 0.0677]

Note: The distributional assumption is that the unobserved γ is $N(\mu,\sigma)$. OPG asymptotic standard errors in column on the right for columns (1) and (2). For columns (2) and (3) the predicted mean of the distribution of γ is calculated at the mean of the covariates. For column (3) the predicted standard deviation of γ is calculated at the mean of the covariates. For regional and seasonal fixed effects we include the p-values of the chi-squared statistics for joint significance. * total number of thresholds including empty intervals (see Table 4).

TABLE 3B

MLE: Unobserved Risk Preferences in the Uninformative Offers Model

CRRA	Dep. Variable: γ		Dep. Variable: γ		Dep. Variable: γ	
	(1)		(2)		(3)	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN						
μ	1.1920	0.0976	1.1900	--	1.1804	--
Individual Covariates:						
Age			0.0158	0.4788	-0.0231	0.4395
$(Age)^2$			-0.0023	0.0112	-0.0004	0.0104
$(Age)^3$			0.0000	0.0001	0.0000	0.0001
Female			-0.0368	0.2465	-0.0300	0.1773
Constant			-41.2089	50.0477	-3.9548	32.0762
Regional F.E. (p-value)			--	[0.63]	--	[0.014]
Income(/10,000)			0.3877	0.7469	-0.0140	0.4615
Time Covariates:						
Time trend			26.0360	29.9447	4.0853	20.4512
Season of the show F.E. (p-value)			--	[012]	--	[0.004]
STANDARD DEVIATION						
σ	1.0486	0.0550	0.9188	0.0780	0.7660	--
Individual Covariates						
Age					1.0461	0.6838
$(Age)^2$					-0.0268	0.0161
$(Age)^3$					0.0002	0.0001
Female					0.5294	0.2928
Constant					-3.6423	16.3282
Income(/10,000)					1.0164	0.7907
Time Covariates						
Time trend					-7.0762	6.5169
Log Likelihood	-168.5295		-149.0711		-138.2643	
No. Thresholds*	517		517		517	
No. Obs.	241		241		241	
No. Left-censored Obs.	131		131		131	
No. Right-censored Obs.	21		21		21	
No. Interval Obs.	89		89		89	
chi-squared [df, p-value]	--		38.9200	[28, 0.0823]	59.7200	[28, 0.0004]

Note: The distributional assumption is that the unobserved γ is N(μ,σ). OPG asymptotic standard errors in column on the right for columns (1) and (2). For columns (2) and (3) the predicted mean of the distribution of γ is calculated at the mean of the covariates. For column (3) the predicted standard deviation of γ is calculated at the mean of the covariates. For regional and seasonal fixed effects we include the p-values of the chi-squared statistics for joint significance. * total number of thresholds including empty intervals (see Table 4).

TABLE 3C

MLE: Unobserved Risk Preferences in the Informative Offers Model

CRRA	Dep. Variable: γ		Dep. Variable: γ		Dep. Variable: γ	
	(1)		(2)		(3)	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN						
μ	1.3688	0.1254	1.3514	--	1.3140	--
Individual Covariates:						
Age			-0.1809	0.5628	-0.8064	0.4436
$(Age)^2$			0.0043	0.0130	0.0192	0.0106
$(Age)^3$			-0.00003	0.0001	-0.0001	0.0001
Female			-0.1239	0.2934	-0.3107	0.1862
Constant			-51.5455	65.7595	-48.8788	27.3729
Regional F.E. (p-value)			--	[0.6293]	--	[0.0303]
Income(/10,000)			0.7785	0.8427	0.3331	0.4227
Time Covariates:						
Time trend			32.9423	39.7442	36.6249	17.5469
Season of the show F.E. (p-value)			--	[0.0748]	--	[0.0028]
STANDARD DEVIATION						
σ	1.2800	0.0842	1.1282	0.1062	0.9773	--
Individual Covariates						
Age					2.1016	0.7522
$(Age)^2$					-0.0498	0.0177
$(Age)^3$					0.0004	0.0001
Female					0.3636	0.2871
Constant					-24.1237	13.3156
Income(/10,000)					1.7102	0.6557
Time Covariates						
Time trend					-4.2717	5.8627
Log Likelihood	-190.1919		-171.0048		-163.8269	
No. Thresholds*	517		517		517	
No. Obs.	233		233		233	
No. Left-censored Obs.	126		126		126	
No. Right-censored Obs.	21		21		21	
No. Interval Obs.	86		86		86	
chi-squared [df, p-value]	--		38.3700	[28, 0.0915]	61.2200	[28, 0.0001]

Note: The distributional assumption is that the unobserved γ is $N(\mu,\sigma)$. OPG asymptotic standard errors in column on the right for columns (1) and (2). For columns (2) and (3) the predicted mean of the distribution of γ is calculated at the mean of the covariates. For column (3) the predicted standard deviation of γ is calculated at the mean of the covariates. For regional and seasonal fixed effects we include the p-values of the chi-squared statistics for joint significance. * total number of thresholds including empty intervals (see Table 4).

TABLE 3D

MLE: Unobserved Risk Preferences in the Static Model in US and Dutch Games

CRRA	Dep. Variable: γ (1)		Dep. Variable: γ (2)		Dep. Variable: γ (3)	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN			**UNITED STATES SAMPLE**			
μ	0.2453	0.0427	0.2546	--	0.2540	--
Individual Covariates:						
Age			-0.0706	0.1701	-0.0421	0.0887
$(\text{Age})^2$			0.0010	0.0043	0.0001	0.0020
$(\text{Age})^3$			-0.000003	0.000035	0.000004	0.000013
Female			-0.2096	0.0896	-0.2290	0.0731
Constant			1.7708	2.2341	1.6123	1.2996
Time Covariates:						
Time trend			-0.0052	0.0026	-0.0081	0.0035
STANDARD DEVIATION						
σ	0.2236	0.0371	0.1796	0.0424	0.1526	--
Individual Covariates						
Age					0.1244	0.1104
$(\text{Age})^2$					-0.0014	0.0012
Female					-0.3585	0.4359
Constant					-4.7064	2.5126
Time Covariates						
Time trend					0.0220	0.0146
Log Likelihood	-52.21		-45.77		-43.03	
No. Obs.	41		41		41	
MEAN			**NETHERLANDS SAMPLE**			
μ	0.3799	0.0488	0.3838	--	0.3793	--
Individual Covariates:						
Age			-0.2477	0.1686	-0.2115	0.1092
$(\text{Age})^2$			0.0057	0.0037	0.0050	0.0025
$(\text{Age})^3$			-0.00004	0.00003	-0.00004	0.00002
Female			0.0132	0.1171	0.0243	0.1002
Constant			3.5671	2.5606	2.9836	1.5676
Time Covariates:						
Time trend			0.0050	0.0038	0.0044	0.0040
STANDARD DEVIATION						
σ	0.2873	0.0463	0.2572	0.0581	0.2580	--
Individual Covariates						
Age					0.0496	0.1050
$(\text{Age})^2$					-0.0006	0.0011
Female					-0.4227	0.3580
Constant					-2.2327	2.4568
Time Covariates						
Time trend					0.0022	0.0140
Log Likelihood	-65.51		-61.33		-60.58	
No. Obs.	44		44		44	

Note: The distributional assumption is that the unobserved γ is N(μ,σ). OPG asymptotic standard errors in column on the right for columns (1) and (2). For columns (2) and (3) the predicted mean of the distribution of γ is calculated at the mean of the covariates. For column (3) the predicted standard deviation of γ is calculated at the mean of the covariates.

TABLE 3E

MLE: Unobserved Risk Preferences in the Uninformative Offers Model in US and Dutch Games

CRRA	Dep. Variable: γ (1)		Dep. Variable: γ (2)		Dep. Variable: γ (3)	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN			**UNITED STATES SAMPLE**			
μ	0.6351	0.1721	0.6281	--	0.5749	--
Individual Covariates:						
Age			0.1031	0.5633	-0.1140	0.3155
$(Age)^2$			-0.0038	0.0139	0.0001	0.0074
$(Age)^3$			0.00004	0.0001	0.0000	0.0001
Female			-0.3567	0.2750	-0.5355	0.2399
Constant			-0.0805	7.4657	3.8539	4.3215
Time Covariates:						
Time trend			0.0022	0.0100	-0.0107	0.0087
STANDARD DEVIATION						
σ	0.8093	0.1194	0.7003	0.1327	0.5194	--
Individual Covariates						
Age					0.1762	0.0910
$(Age)^2$					-0.0016	0.0011
Female					0.1518	0.3775
Constant					-5.2346	1.9792
Time Covariates						
Time trend					0.0201	0.0143
Log Likelihood	-61.12		-55.54		-50.70	
No. Obs.	42		42		42	
MEAN			**NETHERLANDS SAMPLE**			
μ	2.0021	0.2876	2.0860	--	2.0308	--
Individual Covariates:						
Age			-1.9029	1.2025	-2.0270	0.9581
$(Age)^2$			0.0405	0.0262	0.0431	0.0197
$(Age)^3$			-0.0003	0.0002	-0.0003	0.0001
Female			-0.2183	0.6132	0.0223	0.5188
Constant			29.7236	17.9084	31.5350	15.2410
Time Covariates:						
Time trend			0.0303	0.0222	0.0340	0.0186
STANDARD DEVIATION						
σ	1.5072	0.1757	1.3639	0.2587	1.2362	--
Individual Covariates						
Age					-0.0611	0.0962
$(Age)^2$					0.0004	0.0010
Female					-0.1203	0.3358
Constant					1.8538	2.2266
Time Covariates						
Time trend					0.0107	0.0127
Log Likelihood	-64.71		-60.26		-57.76	
No. Obs.	43		43		43	

Note: The distributional assumption is that the unobserved γ is N(μ,σ). OPG asymptotic standard errors in column on the right for columns (1) and (2). For columns (2) and (3) the predicted mean of the distribution of γ is calculated at the mean of the covariates. For column (3) the predicted standard deviation of γ is calculated at the mean of the covariates.

TABLE 4
MLE: Alternative measures of individual wealth

Model:	Static model		Static model		Uninf. Offer Model		Uninf. Offer Model		Inf. Offer Model		Inf. Offer Model	
Wealth measures:	Labor income * 10		Labor income * 0		Labor income * 10		Labor income * 0		Labor income * 10		Labor income * 0	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN μ	3.4130	0.3382	0.3797	0.0196	3.2484	0.0210	0.4635	0.0292	4.0323	0.0272	0.5062	0.0317
STD. DEV. σ	3.4652	0.1764	0.2305	0.0173	0.6013	0.0041	0.3497	0.0284	0.6410	0.0054	0.3583	0.0330
Log Lik.	-203.2		-208.01		-551.56		-191.30		-410.98		-185.53	
No. Obs.	239		236		245		242		232		226	

Note: The distributional assumption is that the unobserved γ is N(μ,σ). OPG asymptotic standard errors in column on the right. The static model required the exclusion of 1 influential observation and the Informative offer model required the exclusion of 3 influential observations. The differences in the number of observations across the different models are due to inconsistent observations, see text for a discussion.

TABLE 5

Consistency of the Choice Set

	Rejection Threshold < Acceptance Threshold (Upper bound < Lower bound)	Rejection Threshold at 0	Large Inconsistencies (Upper bound - Lower bound < -1)	Total Number of Choices in Sample
CRRA		Static Model		
No. Inconsistent choices	22	2	8	*518*
Fraction inconsistencies	0.0426	0.0039	0.0154	--
CRRA		Uninformative Offer Model		
No. Inconsistent choices	16	2	6	*517*
Fraction inconsistencies	0.0309	0.0039	0.0116	--
CRRA		Informative Offer Model		
No. Inconsistent choices	30	10	22	*517*
Fraction inconsistencies	0.0580	0.0193	0.0426	--

Note: Payoff amounts are reported in Euros. The choice set includes decisions on monetary offers only at stages $s > 2$. The total number of stages is 756 (=252 episodes*3 stages). Missing and changes amount to 224 (see Table A1 in Appendix). Of the remaining 532 monetary offers, 13 were made at $s = 5$ with 2 prizes of identical value and are dropped , leaving 519 choices. The threshold-finding routine did not converge in 1 (resp. 2) occasion for the static and uninformative (resp. informative) offers model producing the totals in the last column. Only observations that present inconsistent behavior are dropped. The number of inconsistent episodes is the number of inconsistent choices divided by 2.

TABLE 6A

MLE: Static Model at Large and Small Stakes

CRRA	Dep. Variable: γ All stakes		Dep. Variable: γ Large stakes (Mean)		Dep. Variable: γ Large stakes (Max)		Dep. Variable: γ Small stakes (Mean)		Dep. Variable: γ Small stakes (Min)	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN μ	0.8834	0.0473	**0.9440**	0.04786	**0.9106**	0.04257	**0.5290**	0.2958	**0.4363**	0.5161
Stage 4,5: μ	--	--	**0.97573**	0.06313	**0.9412**	0.0627	--	--	--	--
STD. DEV. σ	0.5341	0.0329	0.29835	0.04556	0.2633	0.0446	1.2618	0.3579	1.7447	0.9631
Stage 4,5: σ	--	--	0.35818	0.07668	0.3258	0.0913	--	--	--	--
Log Likelihood	-177.71		-48.92		-40.21		-24.35		-13.21	
No. Obs.	237		96		91		69		35	
Stage 4,5: log L	--		-37.24		-31.41		--		--	
Stage 4,5: No. Obs	--		65		54		--		--	

Note: The distributional assumption is that the unobserved γ is N(μ,σ). OPG asymptotic standard errors in column on the right. We define 'Large stakes' : 1) a lottery including at least one prize in excess of 250,000 Euros (max option); 2) a lottery with mean payoffs of at least 75,000 Euros (mean option). We define 'Small stakes' : 1) a lottery including no prize in excess of 25,000 Euros (min option); 2) a lottery with mean payoffs of at most 18,000 Euros (mean option). Stage 4,5 indicates that the sample is restricted to those stages.

TABLE 6B

MLE: Uninformative Offer Model at Large and Small Stakes

CRRA	Dep. Variable: γ All stakes		Dep. Variable: γ Large stakes (Mean)		Dep. Variable: γ Large stakes (Max)		Dep. Variable: γ Small stakes (Mean)		Dep. Variable: γ Small stakes (Min)	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN μ	1.1920	0.0976	**1.2679**	0.12951	**1.2480**	0.1367	**0.3985**	0.6171	**0.4867**	0.6696
Stage 4,5: μ	--	--	**1.2169**	0.0935	**1.181**	0.09494	--	--	--	--
STD. DEV. σ	1.0486	0.0550	0.7670	0.1377	0.7856	0.1280	2.4400	0.5608	2.2337	1.2271
Stage 4,5: σ	--	--	0.5486	0.1284	0.5269	0.1472	--	--	--	--
Log Likelihood	-168.53		-50.75		-54.45		-26.36		-13.27	
No. Obs.	241		98		93		70		35	
Stage 4,5: log L	--		-39.72		-34.51		--		--	
Stage 4,5: No. Obs	--		67		56		--		--	

Note: The distributional assumption is that the unobserved γ is N(μ,σ). OPG asymptotic standard errors in column on the right. We define 'Large stakes' : 1) a lottery including at least one prize in excess of 250,000 Euros (max option); 2) a lottery with mean payoffs of at least 75,000 Euros (mean option). We define 'Small stakes' : 1) a lottery including no prize in excess of 25,000 Euros (min option); 2) a lottery with mean payoffs of at most 18,000 Euros (mean option). Stage 4,5 indicates that the sample is restricted to those stages.

TABLE 6C

MLE: Informative Offer Model at Large and Small Stakes

CRRA	Dep. Variable: γ All stakes		Dep. Variable: γ Large stakes (Mean)		Dep. Variable: γ Large stakes (Max)		Dep. Variable: γ Small stakes (Mean)		Dep. Variable: γ Small stakes (Min)	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN μ	1.3688	0.1254	**1.3306**	0.1207	**1.2865**	0.1303	**0.5788**	0.6486	**0.6757**	0.6660
Stage 4,5: μ	--	--	**1.2732**	0.1002	**1.2189**	0.1058	--	--	--	--
STD. DEV. σ	1.2800	0.0842	0.7375	0.1378	0.7585	0.1293	2.2697	0.5128	1.7232	1.0432
Stage 4,5: σ	--	--	0.5874	0.1395	0.5812	0.1474	--	--	--	--
Log Likelihood	-190.19		-53.36		-54.95		-22.97		-8.41	
No. Obs.	233		95		90		68		32	
Stage 4,5: log L	--		-43.65		-38.57		--		--	
Stage 4,5: No. Obs	--		66		55		--		--	

Note: The distributional assumption is that the unobserved γ is N(μ,σ). OPG asymptotic standard errors in column on the right. We define 'Large stakes' : 1) a lottery including at least one prize in excess of 250,000 Euros (max option); 2) a lottery with mean payoffs of at least 75,000 Euros (mean option). We define 'Small stakes' : 1) a lottery including no prize in excess of 25,000 Euros (min option); 2) a lottery with mean payoffs of at most 18,000 Euros (mean option). Stage 4,5 indicates that the sample is restricted to those stages.

TABLE 7

Scores for Large ans Small Stakes episodes reaching stage 5 with moneatry offers at different CRRA parameters

	Small stakes (Min)		Large stakes (Max**)	
	Number of correctly predicted decisions	Fraction of correctly predicted decisions	Number of correctly predicted decisions	Fraction of correctly predicted decisions
γ=0	19	59.38%	9	36.00%
γ=1	23	71.88%	19	76.00%
γ=1.5	22	68.75%	20	80.00%
γ=1.8	22	68.75%	21	84.00%
γ=2	21	65.63%	21	84.00%
γ=3	21	65.63%	20	80.00%
γ=4	20	62.50%	17	68.00%
γ=5	20	62.50%	16	64.00%
*Total**	*32*	100.00%	*25*	100.00%

Note: * For small stakes: 32 of the 78 episodes in Table 2B have both payoffs less or equal to 25,000 Euros. For large stakes: 25 of the 78 episodes in Table 2B have at least a payoff greater or equal to 100,000 Euros. **We change the definition for large stakes from Table 5 in order to make the sample sizes comparable.

TABLE 8

Static Model and Cumulative Prospect Theory.

Weighting:	*Power Function*				*Prelec (1998)*				*Camerer and Ho (1994)*			
CRRA	(1)				(2)				(3)			
	Dep. Variable: γ		Dep. Variable: α		Dep. Variable: γ		Dep. Variable: α		Dep. Variable: γ		Dep. Variable: α	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN												
μ_ι $\iota=\gamma,\alpha$	0.6818	0.2246	0.7219	0.2576	0.4423	0.0594	0.8061	0.2124	0.3970	0.1492	0.6589	0.1514
STD. DEV.												
σ_ι $\iota=\gamma,\alpha$	0.2453	0.3328	0.3146	0.3774	0.1682	0.0755	0.3643	0.3433	0.1567	0.1207	0.2071	0.3057
$\sigma_{\alpha\gamma}$	-0.0669	0.1973	--	--	-0.0429	0.0569	--	--	0.0125	0.0726	--	--
Log L.	-235.43				-241.75				-243.55			
No. Obs.	252				252				252			

Note: The distributional assumption is that the unobserved (γ,α) are joint N(M,Σ). OPG asymptotic standard errors in column on the right for columns (1), (2), and (3).

TABLE 9

Loss aversion and Cumulative Prospect Theory.

Weighting:	*Power Function*				*Prelec (1998)*				*Camerer and Ho (1994)*			
Linear	(1)				(2)				(3)			
	Dep. Variable: λ		Dep. Variable: α		Dep. Variable: λ		Dep. Variable: α		Dep. Variable: λ		Dep. Variable: α	
	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.	Coef.	Std. Err.
MEAN												
μ_ι $\iota=\lambda,\alpha$	3.0806	0.8848	0.9940	0.3185	1.9162	0.4578	1.4068	0.3196	1.9269	0.4611	1.3541	0.3122
STD. DEV.												
σ_ι $\iota=\lambda,\alpha$	1.6392	0.6927	0.1357	1.1613	0.8052	0.4769	0.5755	0.6136	0.7568	0.4473	0.4567	0.6178
$\sigma_{\alpha\lambda}$	0.0577	1.0327	--	--	-0.0937	0.8072	--	--	0.0028	0.7598	--	--
Log L.	-318.53				-260.37				-223.71			
No. Obs.	252				252				252			
Vuong p-va.	0.012				<0.001				<0.001			
Clarke p-va.	<0.001				<0.001				<0.001			

Note: The distributional assumption is that the unobserved (λ,α) are joint N(M,Σ). OPG asymptotic standard errors in column on the right for columns (1), (2), and (3). Vuong (1989) and Clarke (2003) tests p-values are reported against the EUT model of Column (1) Table 3A.

TABLE A1

Stage Choices Tabulation

s =3	Value	Count	Percent	Percent on *m*
	Missing	2	0.79%	--
	Change	24	9.52%	--
	Reject *m*	216	85.71%	95.58%
	Accept *m*	10	3.97%	4.42%
	Tot.	*252*	*100%*	*100%*
s =4	Value	Count	Percent	Percent on *m*
	Missing	12	4.76%	--
	Change	25	9.92%	--
	Reject *m*	147	58.33%	68.37%
	Accept *m*	68	26.98%	31.63%
	Tot.	*252*	*100%*	*100%*
s =5	Value	Count	Percent	Percent on *m*
	Missing	81	32.14%	--
	Change	80	31.75%	--
	Reject *m*	47	18.65%	51.65%
	Accept *m*	44	17.46%	48.35%
	Tot.	*252*	*100%*	*100%*

Note: Reject/Accept refers to monetary offers only. The 252 episodes present a monetary offer 532 times.

TABLE A2

Equilibrium Likelihoods of offers relative to E(Lottery): Informative Offers Model and E(Lottery)>20,000 Euros

Rank		2.00	1.00
Likelihood 5 =	Offer is 0-20%	0.03	*.0189/2*
	Offer is 20-40%	0.15	*.0189/2*
	Offer is 40-60%	0.13	0.09
	Offer is 60-80%	0.44	0.17
	Offer is 80-100%	0.26	0.72

Rank		5.00	4.00	3.00	2.00	1.00
Likelihood 4 =	Offer is 0-20%	0.08	0.07	0.09	0.04	0.03
	Offer is 20-40%	0.60	0.54	0.49	0.63	0.25
	Offer is 40-60%	0.28	0.34	0.35	0.23	0.48
	Offer is 60-80%	*.04/2*	*.0488/2*	*.0702/2*	0.08	0.18
	Offer is 80-100%	*.04/2*	*.0488/2*	*.0702/2*	0.02	0.08

Rank		8.00	7.00	6.00	5.00	4.00	3.00	2.00	1.00
Likelihood 3 =	Offer is 0-20%	*0.01*	0.32	0.39	0.22	0.43	0.35	0.14	0.13
	Offer is 20-40%	*1-.04*	*.68-.03*	0.53	*.7778-.03*	*.575-.03*	0.61	0.79	0.67
	Offer is 40-60%	*0.01*	*0.01*	0.03	*0.01*	*0.01*	*.0323-.02*	*.069-.02*	0.13
	Offer is 60-80%	*0.01*	*0.01*	*.0526-.01*	*0.01*	*0.01*	*0.01*	*0.01*	0.04
	Offer is 80-100%	*0.01*	*0.01*	*0.01*	*0.01*	*0.01*	*0.01*	*0.01*	0.04

Note: The table presents the stage likelihoods of the various offers conditional on the rank of the box held by the contestants. Higher rank means higher prize. The expressions in Italics mean that we spread part of the likelihood over the missing events cells in order to avoid dealing with zero-probability events (indicated in the boxes). This solution does not seem to affect our results. A three-way split of the offers range was also tried and the likelihoods are available from the authors. In expectation the mid-point of the offer relative to the expected values of the lottery is taken. The likelihood above is obtained from the full set of offers to maximize the number of cells but employed only for E(lottery)>20,000 Euros.