Who really wants to be a millionaire? Estimates of risk aversion from gameshow data

(revised version)

Roger Hartley, Gauthier Lanot & Ian Walker

No 747

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Who really wants to be a millionaire?
Estimates of risk aversion from gameshow data*

Roger Hartley, University of Manchester, Manchester M13 9PL, UK

Gauthier Lanot, Keele University, Staffs ST5 5BG, UK

Ian Walker, University of Warwick, Coventry CV4 7AL, UK

21 April 2006

Keywords: Risk aversion, gameshow

JEL Class: D810, C930, C230

Abstract

This paper analyses the behaviour of contestants in one of the most popular TV gameshows ever to estimate risk aversion. This gameshow has a number of features that makes it well suited for our analysis: the format is extremely straightforward, it involves no strategic decision-making, we have a large number of observations, and the prizes are cash and paid immediately, and cover a large range – from £100 up to £1 million. Our data sources have the virtue that we are able to check the representativeness of the gameshow participants. Even though the CRRA model is extremely restrictive we find that a coefficient or relative risk aversion which is close to unity fits the data across a wide range of wealth remarkably well.

Corresponding author: Prof. Ian Walker, Department of Economics, University of Warwick, Coventry CV4 7AL, UK
Email   i.walker@warwick.ac.uk   Fax +44/0 24765 23032   Tel: +44/0 24765 23054

1.   Introduction

The existing empirical literature that addresses the degree of risk aversion is distinguished by both the breadth of its estimates and the sparseness of evidence relating to large gambles. This paper analyses the behaviour of TV gameshow contestants to estimate an EU model. This gameshow has a number of features that makes it well suited for our analysis: the format is straightforward and it involves no strategic decision-making; we have a large number of observations; and the prizes are immediately paid in cash and cover a large range – up to £1 million. We use the data to estimate the degree of risk aversion, and how it varies by the size of the gamble, and by gender. Even though the CRRA model is extremely restrictive we find that it fits the data remarkably well and yields very plausible parameter values. However, we do find that it is rejected in favour of a generalization that allows variable RRA, although the economic significance of the departure from CRRA is small.

Our data comes from what has probably been the most popular TV gameshow of all time, "Who wants to be a millionaire?" (hereafter WWTBAM). Notwithstanding that gameshow data has a number of drawbacks for the purpose of estimating attitudes to risk, this particular game has a number of design features that make it particularly well-suited to our task. In this gameshow the contestant is faced with a sequence of 15 multiple-choice questions. At each stage she can guess the answer to the current question and might double her current winnings but at the risk of losing a question-specific amount, or she can quit and leave the game with her winnings to date. The mechanism of the game is well known and very simple. There is no strategic element, contestants simply play against the house. It is, however, a game where skill matters and this complicates our analysis.

At each stage of the game contestants are reminded that their winnings so far belong to them - to risk, or walk away with. The prizes start at a modest level but, in many countries, reach very high levels. This wide spread of possible outcomes makes WWTBAM a considerable challenge for a simple expected utility CRRA model. The sequential nature of the game gives rise to a further important complication – in all but the last stage of the game, answering a question correctly gives an option to hear the next

question and this itself has a value, over and above the value of the addition to wealth associated with the question. This option value depends on the stage of the game, the contestant's view about the difficulty of subsequent questions, and the degree of risk aversion.

The data was transcribed from the original videotapes of the population of contestants. We established the representativeness of the data by surveying the population of potential contestants (individuals who were invited to appear on each show and from which actual contestants were selected) to obtain information about their characteristics, which we then compare with population surveys such as the Labour Force Surveys.

The paper is structured as follows. In section 2 we outline the existing evidence, including other work that relies on gameshow data. Section 3 explains the operation of the game in more detail. In section 4 we provide a simple model of the game that captures its formal structure so we can show the mechanics of the game in a straightforward way. In section 5 we present the econometric details and the likelihood. In section 6 we give some summary details of the UK data and explain how we estimate risk aversion using this data. In section 7 we present some results and consider possible shortcomings of the work. In section 8 we draw together some conclusions.

2. Existing Evidence

There are several distinct strands to the empirical literature. Firstly, considerable attention has been given to the estimation of Euler equations derived from lifecycle models of consumption and savings (see Hall (1988) and Attanasio and Weber (1989)) where the coefficient on the interest rate in a log-linearised model is the elasticity of substitution. If utility is time separable and exhibits CRRA then this interest rate coefficient is also the inverse of the degree of relative risk aversion, ρ. The typical result in such analyses, usually based on macro data, is that consumption and savings are relatively insensitive to interest rates so the elasticity of intertemporal substitution is small. Thus, the macro-econometric literature largely suggests that the degree of risk aversion is large. Some of this literature[1] considers two assets and backs out risk aversion

---

[1] Notable contributions to this area are Epstein and Zin (1989, 1991).

from the excess returns on equities. Since individual portfolios are typically highly concentrated in relatively safe assets this work implies that the degree of risk aversion is implausibly large. Indeed, the survey of the "equity premium puzzle" by Kocherlakota (1996) suggest estimates of the degree of relative risk aversion that exceed 50[2].

However, this method, which relies on portfolio allocations, has only ever been applied to micro-data in a handful of studies. Attanasio, Banks and Tanner (2002) provide a very plausible estimate ρ of just 1.44 using a large UK sample survey (for the sub-sample at an interior solution i.e. of shareholders).

Jianakopolos and Bernasek (1998) use US survey data on household portfolios of risky assets to examine gender differences. They find that single women are more relatively risk averse than single men - a ρ close to 9 compared to 6. Further differences by age, race, and number of children were also found. Palsson (1996) uses Swedish 1985 cross-section data on portfolios drawn from income tax registers for more than 7,000 households. This study also recognizes the existence of real as well as financial assets and accounts for the gains from diversification that arises when real assets and financial assets are both held. The estimated risk aversion was found to be even higher than Jianakopolos and Bernasek but, in this case, it is not systematically correlated with characteristics apart from finding that risk aversion increases with age.

If utility is intertemporally separable then the extent to which utility varies with income is related not just to consumption and savings, but also to labour supply. This idea has been exploited by Chetty (2003) who derives estimates of risk aversion from evidence on labour supply elasticities. He shows that ρ, in the atemporally separable case, depends on the ratio of income and wage elasticities and that the typical estimates in the labour supply literature implies a ρ of about 1. Indeed, Chetty (2003) shows that under weak separability a positive uncompensated wage elasticity is sufficient to bound ρ to be below 1.25.

A second, albeit small, strand of the empirical literature exploits data on the purchase of insurance cover. Szpiro (1986) is an early example which estimates ρ from time series

---

[2] A number of ideas have been put forward to reconcile the equity premium with estimates of risk aversion obtained by other methods – most plausibly, that the premium is correlated with labour income risk.

data on insurance premia and the amount of domestic insurance cover purchased, and finds ρ to be close to 2. Cicchetti and Dubin (1994) consider a large microdataset on insurance for domestic phone wiring. This paper acknowledges that this insurance is expensive relative to the expected loss (a monthly premium of $0.45 on average compared to just a loss of just $0.26 on average) and yet they found that 57% of customers were enrolled in the insurance scheme. They estimate a hyperbolic absolute risk aversion model and estimate, on average, a rather small degree of ARA. The implied estimate of ρ is of the order of 0.6.

A third, more substantial, strand to the literature takes an experimental approach where participants are offered either real or hypothetical gambles. The best example that uses hypothetical questions is Barsky *et al* (1997) where respondents to the US Health and Retirement Survey were asked if they would accept or reject huge gambles (a 50% chance of doubling lifetime income together with a 50% chance of reducing it by one-fifth/one-third/one-half). Two further distinctive features of this work are that it suggests that there is considerable variation in relative risk aversion, around the mean of about 12, and that relative risk aversion is correlated with risk related behaviour in the data such as smoking, insurance and home ownership.

Donkers *et al* (2001) is a further good example that uses data on preferences over hypothetical lotteries in a large household survey to estimate an index for risk aversion. Their econometric method is semi-parametric, it allows for generalisations of expected utility, and they make weak assumptions about the underlying decision process. They go on to estimate a structural model based on Prospect Theory (see Kahneman and Tversky (1979)). They strongly reject the restrictions implied by expected utility theory and they find that both the value function and the probability weighting function vary significantly with age, income, and the wealth of the individual.

Another example of this strand of the literature is Hartog *et al* (2000) which uses questionnaire evidence on reservation prices for hypothetical lotteries to deduce individual risk aversion. They use three different datasets and find that the mean values of ρ are extremely large (more than 20) in each, which might suggest that the questionnaire method is contaminated by response bias. However recent work by Holt and Laury

(2002) compares estimates from hypothetical lotteries with the same lotteries where the prize is really paid. The authors check whether preferences differ across real and hypothetical lotteries and find that they are similar only for small gambles. The analysis features prizes that range up to several hundreds of dollars which they feel allows them to address the critique raised in Rabin and Thaler (2001) and Rabin (2000). They estimate an expected utility function, using real payoffs, that allows for non-constant RRA. Consistent with Rabin, they find small degrees of RRA (around 0.3) at low prize levels and higher (around 0.9) at high prize levels which, together, fit the data well. However, even the largest payouts considered in Holt and Laury (2002) are small compared to WWTBAM.

The present paper belongs to a final strand of the empirical literature that relies on data generated by gameshow contestants. The earliest example, by Metrick (1993), uses the television gameshow *Jeopardy!* as a natural experiment to estimate a non-linear probit of play that depends on the expected value of the gamble from which the degree of risk aversion can be deduced (other examples are Gertner (1993) and Beetsma and Schotman (2001)). Only small stakes are involved and the implied preferences[3] are not significantly different from risk neutral.

Similarly, Hersch and McDougall (1997) use data from the *Illinois Instant Riches* television gameshow, a high stakes game based on the Illinois State Lottery, to regress the probability of accepting a bet on the bet's expected value and (a proxy for) household income. The estimated structural model is used to infer ρ, and the data again suggests that contestants are nearly risk neutral.

More recently Fullenkamp *et al* (2003) uses the *Hoosier Millionaire* television gameshow to analyze decision-making. Unlike earlier gameshows this involves relatively high stakes. They use a large sample of simple gambling decisions to estimate risk-aversion parameters. However, with this game the prizes are annuities and so their value

---

[3] They also model the ability of contestants to choose strategic best-responses. The results suggest that failure to choose the best-response increases as the complexity of the bet increases. Consistent with much psychological experimental literature, he also finds that the choices that contestants make are affected by the "frame" of the problem.

to contestants will depend on time preference. They find that, assuming a discount rate of 10%, contestants display risk-aversion with the mean ρ ranging from 0.64 to 1.76.

Finally, a new gameshow *Deal or No-deal!* has been analysed in two recent papers – Post *et al* (2006) which considers Australian and Dutch data, and Bombardini and Trebbi (2005) which considers Italian data. This game involves a number of potential prizes, varying from the very small up to very large amounts, which are won essentially by chance. Contestants simply make a sequence of draws to eliminate prizes. On up to nine occasions during this process the contestant is made financial offers to quit and so needs to compare the offer made with the certainty equivalent of remaining prizes. While there is no skill involved contestants need to take a view about what the level of *future* offers to quit might be if they are to appropriately value the option of rejecting a current offer. Unfortunately, it is not clear what the process is that generates offers so this has to be estimated along with the risk aversion parameters. Post *et al* (2006) provide estimates on the degree of relative risk aversion in the range 1 to 2. Using the whole range of data Bombardini and Trebbi (2005) estimate the degree of risk aversion to be unity – i.e. utility is logarithmic, and they go on to estimate degrees of risk aversion with the low stake observations that are indeed close to zero.

3. The WWTBAM Gameshow

WWTBAM has proved to be the most popular TV gameshow ever. The game has been licensed to more than one hundred countries and has been played in more than 60. The structure is well known to contestants who are likely to have watched the experience of many previous contestants. The game features a sequence of fifteen "multiple-choice" questions with associated prizes that, in the UK, carry prizes that start at £100 and (approximately) double with each correctly answered question so that the final question results in overall winnings of £1m. *After* being asked each question the contestant has the choice of quitting with her accumulated winnings or gambling by choosing between the four possible answers given. If the chosen answer is correct the contestants doubles her existing winnings and is asked another question. If the chosen answer is incorrect she gets some "fallback" level and leaves the game. The design is such that the difficulty of questions rise, on average, across the sequence of questions, and the fall back level of

winnings also rises (in two steps). Contestants are endowed with three "lifelines" which are use-once opportunities to improve their odds – so, when faced with a difficult question, contestants may use one or more lifelines to improve their odds. The contestant has the choice between the following lifelines: fifty-fifty ("50:50" hereafter), which removes two of the three incorrect answers; ask-the-audience ("ATA"), which provide the contestant with the distribution of opinion in the audience concerning the correct answer in the form of a histogram; and phone-a-friend ("PAF"), which allows the contestant to call on a friend for a limited amount of time (30 seconds). At least in the UK questions are chosen randomly from a sequence of question banks and there is no attempt to manipulate the fortunes of individual contestants during play.

Contestants are not randomly selected onto the show. The details of how this is done varies across countries but in the UK aspiring contestants ring a premium rate phone number and are asked a question of medium difficulty. If correct, their name is entered into a draw to appear in the studio. Ten names are drawn for each show (plus two reserves). Aspiring contestants can improve their odds of appearing by ringing many times so having many entries in the draw. Once at the recording studio, aspiring contestants compete with each other to provide the fastest correct answer to a single question and the winner is selected to enter the main game.

During play the compère is careful, at least in the UK game, to ensure that contestants are sure they want to commit themselves at every stage – contestants have to utter the trigger phrase "final answer" to indicate commitment. At each of the two fallback stages, the compère hands a cheque to the contestant for that level of winnings and ensures that the contestant understands that this money is now theirs and cannot be subsequently lost. The compère makes every effort to ensure that contestants behave rationally – he strongly discourages quitting at the question corresponding to the fallback levels (where there is no downside risk), and quitting with unused lifelines. He is also careful to ensure that contestants understand the magnitude of the risks that they face during the game. Neither compère, nor the game format, seem to encourage cautious or hesitant contestants to take undue risks.

4. A simple model

4.1. Dynamic aspects of the game

The model of participation we present accounts for the dynamic structure of the game. We focus initially on a simplified version of the game in which contestants are risk neutral and hence are expected income maximisers. We postpone consideration of the "lifelines" to section 4.3 and, for the moment, assume that questions are selected by independent random draws from a pool of questions of identical difficulty.

Let p denote the probability that the contestant (of some given ability) is able to answer correctly a question, where p is a realisation of the random variable P whose cdf is $F : [0,1] \mapsto [0,1]$ (we provide, in the next section, a model for this distribution). Rounds of the game are denoted by the number of questions *remaining*, i.e. $n = N,\ldots\ldots 1$. Let a_n be the accumulated winnings after the contestant has successfully completed $N\text{-}n$ questions and there are n questions remaining. In the televised game $N=15$ and the prizes are given by the sequence $\{a_n\}_{n=1}^{16} = \{1000, 500, 250, 125, 64, 32, 16, 8, 4, 2, 1, 0.5, 0.3, 0.2, 0.1, 0\}$.

Similarly, let b_n be the value of the fallback level of winnings, i.e. the winnings that can be kept in the event of an incorrect answer. In the televised game the sequence of fallback prizes is given by, $\{b_n\}_{n=1}^{15} = \{32, 32, 32, 32, 32, 1, 1, 1, 1, 1, 0, 0, 0, 0, 0\}$.

Now consider the decision problem at the start of the game when the contestant is faced with the first of 15 questions. The value of playing the game is given by $V_{15}(p) = \max\{a_{16}, p(f_{14} - b_{15}) + b_{15}\}$ where $f_{14} = \mathrm{E}[V_{14}(P)]$ is value of continuing optimally and, at this stage $a_{16} = b_{15} = 0$. This is the first stage of a recursion, such that when there are n questions to go and the question asked can be answered successfully with probability p, the value of the game is

$$V_n(p) = \max\{a_{n+1}, p(f_{n-1} - b_n) + b_n\}, \tag{1}$$

where $f_{n-1} = \mathrm{E}[V_{n-1}(P)]$ and we set $f_0 = a_1$. Note that the decision to quit or not to quit is made <u>after</u> the question has been asked.

8

At any round of the game, there exists a critical value of p, $\bar{p}_n = (a_{n+1} - b_n)/(f_{n-1} - b_n)$, such that if $p \le \bar{p}_n$ the individual quits the game and therefore $V_n(p) = a_{n+1}$. Otherwise $p > \bar{p}_n$ and the individual offers an answer to the question and the value of the game is $V_n(p) = p(f_{n-1} - b_n) + b_n$. Note that the immediate value of answering correctly is a_{n+1} and the expected difference, $p(f_{n-1} - b_n)$, represents the "option value" of continuing. These dynamic programming equations lead to the following relationship for $\{f_n\}$:

$$f_{n-1} - f_n = (f_{n-1} - b_n)\int_{\bar{p}_n}^{1} F(p)\,dp . \tag{2}$$

To obtain the likelihood we need to evaluate the probability of winning. The probability of continuing to participate through offering an answer to the n^{th} question, but prior to seeing the next question, is

$$\Pr["\text{Play}"] = 1 - F(\bar{p}_n) \equiv \bar{F}(\bar{p}_n) . \tag{3}$$

The probability of giving a correct answer, having decided to answer, is given by

$$\Pr["\text{Win}"|"\text{Play}"] = \frac{\int_{\bar{p}_n}^{1} p\,dF(p)}{1 - F(\bar{p}_n)} \equiv \frac{\bar{G}(\bar{p}_n)}{\bar{F}(\bar{p}_n)} . \tag{4}$$

Hence the probability of answering correctly is simply $\Pr["\text{Win}"] = \bar{G}(\bar{p}_n)$.

Thus, the likelihood of a contestant reaching round k and then quitting (i.e. refusing to give an answer to question k) is

$$L(k,0) = \{1 - \bar{F}(\bar{p}_k)\} \prod_{n=k+1}^{15} \bar{G}(\bar{p}_n) . \tag{5}$$

The probability of a contestant reaching round k and then giving an incorrect answer is

$$L(k,1) = \{\bar{F}(\bar{p}_k) - \bar{G}(\bar{p}_k)\} \prod_{n=k+1}^{15} \bar{G}(\bar{p}_n) . \tag{6}$$

Finally, the probability of a contestant reaching round 1 and then winning (£1m) is :

$$L(1,.) = \prod_{n=1}^{15} \overline{G}(\overline{p}_n). \tag{7}$$

The model can be adapted easily to allow for risk averse behaviour, indeed prizes simply need to be measured in utility terms, i.e. for some concave increasing utility function $u(x)$, we would have $\{\tilde{a}_n\}_{i=1}^{16} = \{u(a_n)\}_{i=1}^{16}$ and $\{\tilde{b}_n\}_{i=1}^{15} = \{u(b_n)\}_{i=1}^{15}$.

4.2. Questions, Answers and Beliefs

The purpose of this section is to propose a model for the distribution of the beliefs that an individual holds each time she is confronted with a question and a list of possible answers (in the real game, four) of which just one is correct. In this section, and in the next, we take as given that the contestant chooses (if she decided to participate) the answer with the highest subjective probability of being correct. Hence once the distribution of that probability is defined it becomes, in principle, straightforward to describe the probability distribution of the maximum belief and, more generally, of the order statistics.

The process of generating questions (and corresponding answers) that we have in mind can be described as follows. A question, and a list of possible answers, is drawn uniformly (at each stage of the game) from a pool of questions (assuming just one pool for the moment). The question and its possible answers (in a randomized order) are presented to the contestant who is then endowed with a draw from the belief distribution concerning the likelihood of each answer being correct. The formation of beliefs for all contestants is assumed to follow this process in an identical and independent manner. Hence, given a particular question, two otherwise identical individuals can hold distinct beliefs concerning the likelihood of each answer. Furthermore, all individuals are assumed to be able to evaluate the distribution of possible beliefs over the population of questions involved at any given stage of the game.

Formally, suppose that **X** is an n-dimensional random vector on the simplex

$$\Delta_n = \left\{ \mathbf{x} : x_i \geq 0 \; \forall i = 1..n, \sum_{i=1}^{n} x_i = 1 \right\}.$$

10

We assume that $\mathbf{X}$ has the probability density function $\psi_n(\mathbf{x})$ and we require it to exhibit the following symmetry property: if $\mathbf{x}, \tilde{\mathbf{x}} \in \Delta_n$, such that $\tilde{\mathbf{x}}$ is obtained from $\mathbf{x}$ by a permutation of its components, then $\psi_n(\mathbf{x}) = \psi_n(\tilde{\mathbf{x}})$.

For our purposes we may limit our investigation to the cases where $n \le 4$. We construct a family of distributions of beliefs by starting with a probability density function ϕ on $[0,1]$ that is symmetric, i.e. such that $\phi(x) = \phi(1-x)$ for all x in $[0,1]$. Note that ϕ's symmetry implies

$$\int_0^1 \phi(x)(1-x)\,dx = \frac{1}{2} \quad \text{and} \quad \int_0^1 \phi(x)(1-x)^2\,dx = \int_0^1 \phi(1-x)(x)^2\,dx = \mu_2 \quad (8)$$

where μ_2 is the second moment of ϕ. If Φ denotes the distribution function corresponding to ϕ, it is straightforward to show that $\Phi[1-z] = 1 - \Phi[z]$.

Our construction of a class of distributions of beliefs is based on ϕ. In the three cases of interest, we consider the following

$$\psi_2(x_1, x_2) = \frac{1}{2}\left[\phi(x_1) + \phi(x_2)\right], \qquad (9)$$

$$\psi_3(x_1, x_2, x_3) = \frac{1}{3} \sum_{\{i,j,k\} \in \mathscr{P}_3} \phi(x_k)\phi\left(\frac{x_j}{1-x_k}\right), \qquad (10)$$

$$\psi_4(x_1, x_2, x_3, x_4) = \frac{1}{12\mu_2} \sum_{\{i,j,k,l\} \in \mathscr{P}_4} \phi(x_l)\phi\left(\frac{x_k}{1-x_l}\right)\phi\left(\frac{x_j}{1-x_l-x_k}\right), \qquad (11)$$

where, for any n, $\mathscr{P}_n$ is the set of all permutations of $\{1,...,n\}$, and $\mu_2 = \int_0^1 x^2\phi(x)\,dx$.

In each case the role of the summation of the set of permutations arises because of the unobserved random (uniform) order in which the contestant answers are presented to the participant. Because ϕ is itself symmetric and $\sum_{i=1}^n x_i = 1$, some (more or less obvious) simplifications are possible, we have

$$\psi_2(x_1, x_2) = \phi(x_1), \qquad (12)$$

11

$$\psi_3\left(x_1, x_2, x_3\right) = \frac{2}{3}\left[\phi\left(x_2\right)\phi\left(\frac{x_2}{1-x_1}\right) + \phi\left(x_2\right)\phi\left(\frac{x_3}{1-x_2}\right) + \phi\left(x_3\right)\phi\left(\frac{x_1}{1-x_3}\right)\right], \quad (13)$$

$$\psi_4\left(x_1, x_2, x_3, x_4\right) = \frac{1}{6\mu_2} \sum_{\substack{l,k=1,\ldots,4 \\ k \neq l}} \phi\left(x_l\right)\phi\left(\frac{x_k}{1-x_l}\right)\phi\left(\frac{x_j}{1-x_l-x_k}\right), . \quad (14)$$

These simplifications are useful in practice since the number of terms involved is halved. Note that in each case it can be verified that the integral of ψ_n over Δ_n is unity, and that ψ_n satisfies the symmetry property required above. In all cases, if ϕ is the density of the uniform distribution between 0 and 1, then ψ_n is the uniform distribution over Δ_n.

This specification of the beliefs distribution is of course restrictive even among the distributions satisfying the imposed symmetry property. It leads, however, to simple specifications for the distribution of the order statistics and for the distribution of the maximum amongst $\left(x_1, \ldots, x_n\right)$. See appendices A1 and A.4.

4.3. Lifelines

Contestants are endowed with three lifelines, described earlier, that can each be played once at any time in the game. Two, or even all three, lifelines may be used on one question. Let us first consider the game with only one remaining lifeline. To clarify the difference the lifeline makes, Figure 1 presents the decision trees at stage n without or without the lifeline. To account for the lifeline, the state space clearly has had to be extended. We write $W_n\left(\mathbf{p};\gamma\right)$ for the expected value of the game to a contestant faced with a question with belief vector $\mathbf{p}$, when $\gamma = 0$ if the lifeline has been used and $\gamma = 1$ if it is still available. Whether to use a lifeline or not may depend on all components of $\mathbf{p}$ so the value is a function of the whole vector of subjective probabilities. However, $p = \max_i p_i$ is a sufficient statistic for $\mathbf{p}$ in the contestant's decision problem with no lifeline left and we will write this value function as $V_n\left(p;0\right)$.

In what follows we assume that the lifeline amounts to a draw of a new belief, say $\mathbf{q}$, given $\mathbf{p}$, the current belief. For example, the use of 50:50 reduces two components of the belief vector to 0. For the other lifelines the audience and/or one (among several)

friends will provide some information which is then combined with the initial belief $\mathbf{p}$. The new belief is the outcome of this process and $\mathbf{q}$ is then used instead of $\mathbf{p}$ in the decision problem. We therefore assume that the conditional distribution function of $\mathbf{q}$ given $\mathbf{p}$ is well defined for each lifeline. Finally we define

$$k_n(\mathbf{p}) \equiv E_{\mathbf{q|p}}\left[V_n(q,0)|\mathbf{p}\right], \tag{15}$$

to be the value of playing the lifeline at stage n where $q = \max_i q_i$.

The values $W_n(\mathbf{p},1)$ and $V_n(p,0)$ are then related according to the updated dynamic programming equations below. When no lifeline is left we have the familiar equation:

$$V_n(p,0) = \max\left\{a_{n+1}, p\left(f_{n-1}(0) - b_n\right) + b_n\right\}, \tag{16}$$

where $f_n(0) = E\left[V_n(P;0)\right]$. When the lifeline remains the contestant will choose the largest of the three options in the first choice line in Figure 1b (quit, use the lifeline, or play), where :

$$W_n(\mathbf{p},1) = \max\left\{a_{n+1}, k_n(\mathbf{p}), p\left(f_{n-1}(1) - b_n\right) + b_n\right\}, \tag{17}$$

and $f_{n-1}(1) \equiv E\left[W_{n-1}(\mathbf{P},1)\right]$ and $f_0(1) = f_0(0) = a_1$.

Note that contestants will never strictly prefer to quit with a lifeline left unused. However, it is still possible that, for some $\mathbf{p}$, a contestant may be indifferent between quitting and using the lifeline if she would subsequently choose to quit for any realisation of $\mathbf{q}$ contingent on $\mathbf{p}$. For example, a contestant for whom $\mathbf{p} = (\frac{1}{4},\ \frac{1}{4},\ \frac{1}{4},\ \frac{1}{4})$ will have $q=1/2$ no matter which two incorrect answers are removed if she were to play "50:50" (which two answers are discarded is decided when the question is constructed). This may fall short of the initial value of answering the question. A contestant who would reject such a "50:50" gamble would place no value on the lifeline. Except in these circumstances, the lifeline will be used if $a_{n+1} \leq k_n(\mathbf{p})$, otherwise the contestant will answer and retain the lifeline for future use.

The treatment of contestants with more than one remaining lifeline is a straightforward extension of this approach. Details are given in appendix A.2.

5. Econometric specification and estimation

5.1. Specification of the belief distribution

The distribution of the beliefs is one of the main elements of the model since it describes the distribution of the unobservables in the model. Under the assumptions we make below (see section 5.2) the joint density $\psi_4(\ )$ can be constructed from some symmetric density ϕ over the unit interval. We assume that $\phi(x)$ is the density of a symmetric Beta random variable, $B(\alpha,\alpha)$ on [0,1], i.e.:

$$\phi(x) = \frac{\Gamma(2\alpha)}{\Gamma(\alpha)^2} x^{\alpha-1}(1-x)^{\alpha-1}, \tag{18}$$

with α some positive parameter, and where $\Gamma(.)$ is the gamma function. Any random variable distributed according to $B(\alpha,\alpha)$ has expectation ½ and second moment $\mu_2 = \frac{1}{2}\left(\frac{\alpha+1}{2\alpha+1}\right)$.

In what follows, it will prove convenient to use draws from the joint distribution of ordered statistics of the belief distribution. Because of the symmetry assumptions that we impose on $\psi_4(\ )$, the joint density function of the order statistics (i.e. the vector of beliefs in decreasing order) is simply $4!\psi_4(\tilde{\mathbf{p}})$, where $\tilde{\mathbf{p}}$ is a vector of values ordered in decreasing order.

Since it is straightforward to rank four numbers in decreasing order, the only remaining issue is to draw from a multidimensional random variable with joint density $\chi_4(\ )$. In Appendix A.3 we show how this can be achieved using three independent draws from Beta distributions with parameters $(\alpha,\alpha+2),(\alpha,\alpha+1)$ and (α,α), respectively[4].

---

[4] We assume that the value of the parameter of the Beta distribution which generates the beliefs at each stage of the game varies with the number of questions left in the following manner:

14

5.2. Likelihood

The contribution of an individual history to the likelihood is the product of the probabilities of success, and of the particular pattern of use for the lifelines for that individual history, up to and including the penultimate question, multiplied by the probability that for her last question she wins a million, loses or quits and the observed use of the lifelines for this last question.

We assume that the expected utility function takes the CRRA form $U(c) = (c+\gamma)^{1-\rho} / (1-\rho)$ where initial wealth, γ, is treated as a parameter to be estimated.

Hence the contribution of contestant i's history, which ends at stage n_i^*, to the likelihood has the general form

$$\hat{L}_{S,i}\left(\left\{\left(LL(k,i), ll(k,i)\right)\right\}_{k=1}^{n_i^*}; (\alpha,\rho,\gamma,\kappa); \left(\hat{\lambda},\hat{\nu}\right)\right) = \left\{\prod_{k=1}^{n_i^*-1} \hat{\Omega}_{ll(k,i),k}^{LL(k,i)}\right\} \hat{\Omega}_{ll(n_i^*,i),n_i^*}^{LL(n_i^*,i)}, \qquad (19)$$

where $LL(k,i)$ indicates the number and nature of the lifelines available to the contestant i at stage k, and $ll(k,i)$ selects the relevant probability depending on the lifeline used by contestant i at stage k. Hence $(\alpha,\rho,\gamma,\kappa)$ is the vector of parameters of interest, i.e. α is the parameter of the belief distribution, ρ is the coefficient of relative risk aversion, γ is a scaling factor in the utility function, and κ is the probability that the friend knows the correct answer when the PAF lifeline is used[5]. Finally $\left(\hat{\lambda},\hat{\nu}\right)$ are the independent estimates of the parameters of the density of the updated belief which results from the use of "Ask the Audience" – see Appendix A2.3 for further details.

$$0.1\left(\frac{n}{10}\right)^3 + \delta_1\left(1-\left(\frac{n}{10}\right)^3\right) + \delta_2\left(\frac{n}{10}\right)\left(1-\frac{n}{10}\right),$$

where δ_1 and δ_2 are two additional parameters we estimate. This specification allows for a relatively flexible association between the stage of the game and the distribution of beliefs at the cost of a small number of additional parameters. For example if $\delta_1 = 1.25$ and $\delta_2 = -1.725$ the relationship is decreasing and close to linear, if instead $\delta_2 = -3.427$ the relationship is decreasing and convex , and decreasing and concave if $\delta_2 = 0$.

[5] In the absence of any reliable information on the quality of the friend's opinion, we allow the data to decide on the value of κ.

6. Data

The operator, Celador PLC, selected 10 names at random from a (large) list of entrants, for each show broadcasted, who had successfully answered a simple screening question over a premium rate phone line. These 10 individuals attended the recording session for their show where they would compete against each other to be quickest to correctly answer a general knowledge question in a simple first round game known as the "Fastest Finger". The winner of this initial round then competes, against the house, in the second round sequence of multiple choice questions. Typically each show would have time for two or three second round contestants. Contestants still playing at the end of the show would continue at the start of the next show.

Our data comes from two sources. We have data that has been extracted from videotapes of the broadcast shows, kindly made available to us by Celador. These tapes cover all shows in the eleven series from its inception in September 1998 to June 2003. This gives us information on the behaviour of 515 contestants[6] who played the second round sequence of multiple choice questions. However, a major concern about the findings of the gameshow literature is that the data is generated by selected samples[7]. To investigate this issue a questionnaire was sent to all of the 2374 potential contestants (except one) who had ever been invited to the studio for all UK shows in the first eleven series of shows broadcast. The questionnaire was designed to identify differences between contestants and the population as a whole. The questions aimed to provide data that was comparable to that available from official social surveys of large random samples of the population[8].

---

[6] We drop the shows that featured couples (including twins, father/sons, professors/freshers) and celebrities. One show, where a contestant was the subject of litigation, was not available to us. We also dropped the single observation who quit below £1000.

[7] In fact, Hersch and McDougall (1977) and Fullenkamp *et al* (2003) do report some comparisons between contestants and the population and find no significant differences on observable characteristics except for lottery play. This latter difference is unsurprising since all contestants have had to have played the lottery and won in order to appear on these shows. In the UK, lottery contestants do seem to have different characteristics than non-contestants (see Farrell and Walker, 1999).

[8] To protect confidentiality, we were not able to match the questionnaire data to the gameshow videotape information so we ensured that the questionnaire also contained information about play during the game.

Questionnaire replies were received in 791 cases, a response rate of 33%, where 243 (32%) of these cases were Fastest Finger winners and so played the second round game. These 243 represent a response rate of 47% of the population of second round contestants. Not surprisingly, these second round contestants were more likely to respond to the survey because they were well disposed towards Camelot, having had the opportunity to win considerable amounts of money. It was immediately obvious that men were heavily overrepresented in both datasets. Table 1 shows the means of the data for the second round competitors and for the non-competitors. The "Fastest Finger" winner who go on to become WWTBAM competitors are more likely to be male, are a little younger, and have slightly longer education than those that failed at this first round. The rest of the table shows the corresponding information from various social surveys, re-weighted to match the gender mix in the questionnaire data.

Once the population datasets are re-weighted the observable differences between the questionnaire data and the population survey data tend to be quite small. Two variables are particularly worthy of note: the proportion of individuals who report that their household's contents are not insured is similar to the population value (in fact slightly smaller suggesting more risk aversion); and the proportion who report being regular lottery ticket purchasers is also quite similar. Thus, our questionnaire dataset does not suggest that those that play (in the second round of) WWTBAM are heavily selected according to observable variables – except gender. Indeed, for those variables which might be expected to reflect risk attitudes we find no significant differences with our population surveys.

However, whether the same can be said about the videotape information which is the population of WWTBAM contestants depends on the questionnaire respondents being representative of this underlying population. Thus, in Table 2, we compare the questionnaire data for the sample of 243 contestants who responded to the questionnaire with the population of 515 actual contestants. We have no consistent information on the characteristics of contestants in the population apart from that which is recorded on the videotapes. Thus, Table 2 records only gender and the outcomes of play. There are no significant differences in gender and although the outcomes information shows, as might be expected, that the questionnaire respondents were bigger winners on average, these

17

differences are not significant. Thus, we can have some confidence that the representativeness of contestants (in the questionnaire data) carries over to the population data in the videotapes.

The distribution of winnings, for the second round contestants, depends on whether the contestant quit or failed to answer the last question asked. Almost all contestants who survived beyond £125,000 quit rather than failed – only one contestant failed at £500,000 and so went away with just £32,000 instead of quitting and going away with £250,000. Only three contestants failed at a sub £1000 question and went away with nothing. Three contestants won the £1m prize. Two-thirds of contestants quit and one-third failed. "Failures" left the studio with an average of £17,438 (£15,000 for women and £18,244 for men) while "quitters" went away with an average of £72,247 (£68,182 for women and £73,411 for men)[9].

Figure 2 presents a scatter plot of the value of the last question seen against the amount actually won. The scales for each axis are logarithmic (i.e. a minor tick on the axis indicates an integer multiple of the major tick to the left or below), and the data points have been jittered slightly to give an impression of the density around each point. The off-diagonal winnings at £1000 and £32000 stand out since these are the only amounts that can be reached from questions with a higher value when the contestant decides to play and loses. The on-diagonal data points represent individuals who have decided to quit when facing a question which they feel is "too difficult".

Figure 3 shows the overall distribution (the continuous line) of fails/quits as the game progresses from the first relevant question (i.e. when only 10 questions are left to play before winning the million pounds) to the last. Note that there are a disproportionate number of male contestants and we have represented the same proportions for each gender. Men tend to fail less in earlier stages of the game, while women tend to quit earlier. Although in some cases the differences are large, the overall pattern is comparable across genders.

Finally, the use of lifelines is an important part of observed behaviour that our model attempts to explain. Out of the 515 contestants, 501 played ATA, 488 played

---

[9] Here we categorise those that won the maximum £1m as quitters.

50:50, and 484 played PAF. There were no examples of individuals quitting with unused lifelines. There was a systematic tendency for lifelines to be played in order. ATA was played, on average, with 8.5 questions remaining; 50:50 was played with, on average, 7.0 questions left; and PAF was used with just 6.9 questions remaining, on average.

7. Estimation and results

The estimation of the parameters of the model requires that we first estimate our model for the histograms that the lifeline "Ask the Audience" produces. This histogram records the percentage of the audience that selects each of the four possible answers. This preliminary estimation is used to determine the parameters $\left(\hat{\lambda}, \hat{\nu}\right)$ of the distribution described in Section A.2.3 of the appendix, where further details are given.

To do so we use the data we have collected on these histograms and on our knowledge of the correct answer. Section A.2.3. in the appendix describes formally this aspect of the model. These parameter estimates allow us to evaluate the quality of the lifeline "Ask the Audience". The parameter estimates are presented in Table 3. Assume that the first contestant answer is the correct one, these estimates imply that on average we expect the lifeline "Ask the Audience" to produce the histogram 0.63, 0.12, 0.12, 0.12, (i.e. $\left(\dfrac{\hat{\lambda}}{\hat{\lambda}+3\hat{\nu}}, \dfrac{\hat{\nu}}{\hat{\lambda}+3\hat{\nu}}, \dfrac{\hat{\nu}}{\hat{\lambda}+3\hat{\nu}}, \dfrac{\hat{\nu}}{\hat{\lambda}+3\hat{\nu}}\right)$). Treating these parameters as constants we then proceed to estimate the remaining parameters of the model.

Table 4 presents the estimates of the preference parameters for two different specifications of the utility function and for the whole sample as well as for the samples of men and women separately. In addition to our preference parameters, Table 4 gives an estimate of the probability that the chosen friend, in the case of the PAF lifeline, knows the correct answer which, since we have no reliable data on this[10], is estimated jointly with the preference parameters.

The first column of Table 4 gives the parameter estimates of the baseline CRRA model for the pooled sample of men and women as a whole and assumes that the scale

---

[10] Many friends prove not very helpful and do not even hazard a guess.

parameter, γ, (which one might think of as initial wealth) is constant across individuals. Our estimated value for the coefficient of relative risk aversion (first column) is remarkably close to 1 (although statistically significantly different from 1). The parameter γ, which might be interpreted as reference wealth measured in thousands of pounds, is significantly estimated at 0.41.

The remainder of the table relaxes the restrictions that: risk aversion is constant, men and women are the same; and that γ is a constant. In particular, we present results for the generalization of the CRRA and the CARA utility functions known as the Hybrid Power-Expo (see for example, Holt and Laury, 2002). This specification depends on three parameters (instead of two for CRRA) in the following way:

$$u(x) = \frac{-\mathrm{sign}(\alpha)}{\alpha} \exp\left(-\frac{\alpha}{1-\rho}(x+\gamma)^{1-\rho}\right).$$

It naturally nests both the CRRA (in the limit when $\alpha \to 0$) and the CARA (when $\rho = 0$) specification and is therefore a convenient alternative specification to consider.

We provide parameter estimates of this specification for the pooled sample and for men and women separately. In the pooled sample case, although α is statistically significant, the other parameters change little. In the estimation for the separate male and female samples, it is clear that the estimates do not exhibit large differences between men and women. The final three columns of the table presents estimates for this hybrid specification but allowing for some unobserved heterogeneity in γ. We assume that γ is independently and identically distributed across individuals and, for the purpose of the estimation, we endow γ with a discrete (2 mass point) distribution. This clearly rejects the one mass point assumption but again the preference parameters are effectively identical.

The estimation of the alternative specification over the whole sample or over the samples of men and women separately gives different parameter estimates although these differences are again not very large. Even allowing for unobserved heterogeneity does not modify our findings dramatically. We interpret these findings as providing support for the robustness of the estimates based on the CRRA specification.

Two additional parameters which allow for the distribution of the initial belief to change with each round of the game are also estimated (not shown). To illustrate how the distribution of the beliefs changes as the game progresses we have calculated, in Figure 4, the distribution of the maximum belief when 1 question, and 3, 5, 8, and 10 questions, remain to be played. Note that the questions get more difficult in the sense that the distribution at any stage is stochastically dominated by the distribution at earlier stages. Indeed, with ten questions to go, the corresponding probability density function (i.e. the slope of the distribution function) increases up to, and is concentrated close to, $p=1$. Whereas, with few questions left the probability density function peaks before 1, and is concentrated around $p=0.45$ for the last question.

Figure 5 shows the value of playing the game, V, implied by the CRRA estimates for the pooled sample, as a function of the number of questions remaining (on the horizontal axis) and the number and nature of the lifelines left. As we would expect the value of playing rises as the number of remaining questions falls and lifelines add positive value to playing. ATA appears to be the most valuable lifeline while 5050 and PAF have almost identical values. In fact, the model predicts that the lifeline ATA is almost as valuable as having both 5050 and PAF together.

In Figure 6 we show the predicted probabilities of quitting and failing at each question, computed from the CRRA specification estimated over the whole sample, and compare these with the observed distributions. There are many fails and no quits when there are four more questions to come – i.e. when confronted with the £64,000 question – since there is no risk at this point. We broadly capture the peak in quits immediately before this point but underestimate the number immediately afterwards. We overpredict fails for very easy questions, while we underpredict quits.

Figure 7 and 8 illustrate the differences between three possible specification of the utility function: $u(x) = \ln(x + \gamma)$, the estimated CRRA, and the estimated hybrid power expo specification. The scale and the location of each utility function are normalised in Figure 7 such that the comparison in terms of curvature is meaningful. The horizontal axis uses a log scale so that the $\rho=1$ case appears linear. Differences in curvature exist although, from inspection of the figure, these differences are quite small. Figure 8 shows

the coefficient of relative risk aversion (in terms of $x + \gamma$) for the CRRA and the hybrid specification. In keeping with our observation of figure 7 the gradient is modest – it seems that a CRRA of unity is adequate to explain behaviour over a wide range[11].

Finally, in Table 5, we compute the certainty equivalents of the gambles at each stage of the game for three different types of situations (based again on the CRRA specification estimated on the whole sample). The certainty equivalent measures, in money terms, the value of being able to play the game and take into account the option value of being able to play further if the contestant is successful at the current stage. Moreover we present similar calculations for the value of playing the lifeline (again given a particular draw).

The top third of the table corresponds to a situation where the contestant feels she is confident that she knows the answer (i.e. with a belief draw such that she feels that the first possible answer is the correct one with probability 0.9). The middle third of the table corresponds to a situation where the contestant is less confident about the correct answer (i.e. she feels that the first possible answer is the likely correct answer but only with probability 0.6). While the bottom third corresponds to a situation where the contestant is quite unsure about the correct answer (i.e. she feels that the correct answer is the first, but without much confidence).

The table should not be read as a sequence of events for three different types of contestant, rather it lists the possible situations which any contestant might find themselves in. For example individuals are unlikely to find themselves unsure or certain about every single question, rather an individual might be confident about the £16000 question and so have a certainty equivalent of £47120 at that stage, and then be unsure about the £32000 question so that the certainty equivalent of playing has only increases to £50560.

Table 5 also shows the value associated with retaining the lifeline for future use. For example a confident candidate for the £32000 question, has a certainty equivalent of £89860 if she does not play the "50:50" lifeline, compared to a certainty equivalent of

---

£73340 if she decides to use it at this stage. That is, the value of exercising the option of playing the lifeline at this stage is negative because she is confident about the answer to this question and may not be so confident about future questions. Imagine, furthermore, a contestant who has reached the £125000 stage and is confident about the answer, then playing the lifeline reduces the certainty equivalent of being in that position. If on the other hand, in the unlikely event that, she reaches the £250000 stage and is equally confident about knowing the correct answer to the question, then the situation is reversed: she should use her lifeline since the certainty equivalent is larger.

Clearly the belief has a substantial effect on the certainty equivalents. Indeed our model predicts that if faced with either of the second ($p=0.6$) or third ($p=0.4$) belief draws, contestants would be prepared to pay sizeable amounts (amounts larger than £300,000 in the case of the million pound question) to avoid having to answer the question. If EU is logarithmic, contestants will only attempt the £1 million question if they are more than 86% confident.

8. Conclusions

This paper provides new evidence about the degree of individual risk aversion. The analysis is firmly embedded in the expected utility paradigm. Our contribution is to the exploit a "field experiment", based on the popular gameshow, "Who Wants to be a Millionaire?", where the outcomes can vary enormously across contestants, but in a known fashion. This provides a check on laboratory experiments where, although no skill is involved to complicate the analysis, the expected value of the outcome in such experiments is typically small. The range of possible outcomes in WWTBAM is substantially larger than for most of the existing gameshow evidence. The downside of our data source is that we need to make parametric assumptions concerning the distribution of prizes through our assumptions concerning the distribution of beliefs. This is usually not required in a laboratory experiment or in field experiments where the outcomes are determined only by statistical chance. However, we feel that we are in a better position to measure the coefficient of risk aversion than most researchers who use observational data sources where the range of outcomes and their distribution is also usually unknown and need to be inferred.

We also use our data to estimate the value of additional information to contestants in this game of skill. Finally, we were also able to get detailed questionnaire information that provide some reassurance as to the representativeness of contestants. Once we rebalanced the data to account for the disproportionate number of male contestants, our data appears to be representative of the UK population, both in terms of observable characteristics and in terms of other aspects of risk-taking behaviour.

Perhaps surprisingly, we find this parsimonious model is broadly effective in explaining behaviour in this simple game. Our headline result, that the degree of CRRA is approximately 1 with a high degree of precision, is consistent with the results of recent work on the "Hooster Millionaire" and "Deal, No deal" gameshows which are the only other games which feature, like WWTBAM, such large stakes.

References

Adda, J. and R. Cooper (2003), Dynamic Economics, MIT Press: Cambridge, MA.

Attanasio, O., J. Banks and S. Tanner (2002), "Asset Holding and Consumption Volatility", Journal of Political Economy, 110(4), 771-92

Attanasio, O. and G. Weber (1989), "Intertemporal Substitution, Risk Aversion and the Euler Equation for Consumption", Economic Journal, 99, 59-73.

Bombardini, M. and F. Trebbi (2005), "Risk aversion and expected utility theory: a field experiment with large and small stakes", mimeo, Harvard.

Barsky, R. B., Juster, F. T., Kimball, M. S. and Shapiro, M. D. (1997). 'Preference Parameters and Behavioural Heterogeneity: an Experimental Approach in the Health and Retirement Study', Quarterly Journal of Economics, 112(2), 729–58.

Beetsma, R.M. and P.C. Schotman (2001), "Measuring Risk Attitudes in a Natural Experiment: Data from the Television Game Show Lingo", Economic Journal, 111(474), 821-48.

Bombardini, M. and F., Trebbi (2005), "Risk Aversion and Expected Utility Theory: A Field Experiment with Large and Small Stakes", *mimeo* Harvard University.

Chetty, R. (2003), "A New Method of Estimating Risk Aversion", NBER WP 9988.

Ciccheti, C.J. and J.A. Dubin (1994), "Microeconometric Analysis of Risk Aversion and the Decision to Self-Insure", Journal of Political Economy, 102(1), 169-86.

Davidson and McKinnon (2004)

Didonato, A. and A.H. Morris (1986), "Computation of the Incomplete Gamma Function Ratios and their Inverse", ACM Transactions on Mathematical Software, 12, 377-393.

Donkers, B., B. Melenberg and A. van Soest (2001), "Estimating risk attitudes using lotteries - a large sample approach", Journal of Risk and Uncertainty, 22, 165-95.

Epstein, L.G. and S.E. Zin (1989), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework", Econometrica, 57(4), 937-69.

Epstein, L.G. and S.E. Zin (1991), "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis", Journal of Political Economy, 99(2), 263-86.

Farrell, L. and I. Walker (1999), "The Welfare Effects of Lotto: Evidence from the UK", Journal of Public Economics, 72, 99-120.

French, S. (1980), "Updating of Belief in the Light of Someone Else's Opinion", Journal of the Royal Statistical Society, 143, 43-48.

Fullenkamp, C., R. Tenorio and R. Battalio (2003), "Assessing Individual Risk Attitudes using Field Data from Lottery Games", Review of Economics and Statistics, 85(1), 218-26.

Gertner, R. (1993), "Game Shows and Economic Behavior: Risk-Taking on Card Sharks", Quarterly Journal of Economics, 108(2), 507-21.

Gouriéroux, C. and A. Monfort (1996), Simulation-Based Econometric Methods, Oxford University Press: Oxford.

Hall, R.E. (1988), "Intertemporal Substitution in Consumption", Journal of Political Economy, 96(2), 339-57.

Hartog, J., A. Ferrer-i-Carbonell and N. Jonker (2000), "On a Simple Measure of Individual Risk Aversion", Tinbergen Institute Discussion Paper TI 2000-074/3.

Hersch, P.L. and G.S. McDougall (1997), "Decision Making under Uncertainty When the Stakes Are High: Evidence from a Lottery Game Show", Southern Economic Journal, 64(1), 75-84.

Holt, C.A. and S.K. Laury (2002), "Risk Aversion and Incentive Effects", American Economic Review, 92(5): 1644-55.

Jianakoplos, N.A.. and A. Bernasek (1998), "Are Women More Risk Averse?", Economic Inquiry, 36(4), 620-30.

Kahneman, D. and A. Tversky (1979), "Prospect theory: An analysis of decision under risk", Econometrica, 47, 263-291.

Kocherlakota, N.R. (1996), "The Equity Premium: It's Still a Puzzle", Journal of Economic Literature, 34(1), 42-71.

Metrick, N. (1993), "A Natural Experiment in Jeopardy!", American Economic Review, 85(1), 240-53.

Palsson, A-M. (1996), "Does the Degree of Relative Risk Aversion Vary with Household Characteristics?", Journal of Economic Psychology, 17(6), 771-87.

Poirer, D. (1995) Intermediate Statistics and Econometrics: A Comparative Approach, Cambridge, MA: MIT Press.

Post, T., G. Baltussen, and M., Van den Assem (2006): "Deal or no Deal: Decision making under risk in a large payoff gameshow", Tinbergen Institute WP 009/2.

Rabin, M. (2000), "Risk Aversion and Expected-Utility Theory: A Calibration Theorem", Econometrica, 68(5), 1281-92.

Rabin, M. and R.H. Thaler (2001), "Anomalies: Risk Aversion", Journal of Economic Perspectives, 15(1), 219-32.

Szpiro, G.G. (1986), "Measuring Risk Aversion: An Alternative Approach," The Review of Economics and Statistics , 68(1), 156-159.

Train, K.E. (2003) Discrete Choice Methods with Simulation, Cambridge University Press: Cambridge.

Appendix

A. 1 Distribution of the maximum belief

The dynamic model outlined involves the distribution, F, of the contestant's assessment of her chance of answering the question successfully. Without lifelines, $F \equiv F_n$ is the distribution of $\max_{\mathbf{x} \in \Delta_n}(\mathbf{x})$ if $\mathbf{x}$ has the probability density function $\psi_n(\mathbf{x})$. Indeed $\max_{\mathbf{x} \in \Delta_n}(\mathbf{x})$ measures the individual assessment of her likelihood of answering the question correctly when faced with n alternatives. In this section, we describe formulae for the distribution of the highest order statistic:

$$F_n(z) \equiv \Pr\left[\bigcap_{i=1}^{n}\{X_i < z\}\right], \tag{20}$$

given that $\mathbf{X}$ is distributed with density function ψ_n. In particular we can show that

$$F_2(z) = \begin{cases} 0 & \text{if } z \leq \frac{1}{2} \\ 2\Phi[z] - 1 & \text{if } \frac{1}{2} \leq z \leq 1, \\ 1 & \text{otherwise.} \end{cases} \tag{21}$$

and, as a consequence, the density function $f_2(z)$ has support $\left[\frac{1}{2}, 1\right]$ where is satisfies

$$f_2(z) = 2\phi[z]. \tag{22}$$

The distribution function at higher orders can be obtained from F_2 recursively. Whenever $z \in (0,1)$, we have

$$F_3(z) = 2\int_{1-z}^{1} F_2\left(\frac{z}{y}\right) y\phi(y)\, dy, \tag{23}$$

$$F_4(z) = \frac{1}{\mu_2}\int_{1-z}^{1} F_3\left(\frac{z}{y}\right) y^2\phi(y)\, dy, \tag{24}$$

and the relevant density functions, say f_3 and f_4, can be shown to exist and to be continuous everywhere inside $(0,1)$. For example, in the uniform case where $\phi(x)=1$ for $x \in [0,1]$, and 0 elsewhere, we find that

$$F_4(z) = \begin{cases} 0 & \text{if } 0 \leq z \leq \frac{1}{4} \\ (4z-1)^3 & \text{if } \frac{1}{4} \leq z \leq \frac{1}{3} \\ -44z^3 + 60z^2 - 24z + 3 & \text{if } \frac{1}{3} \leq z \leq \frac{1}{2} \\ 1 - 4(1-z)^3 & \text{if } \frac{1}{2} \leq z \leq 1 \end{cases} \tag{25}$$

In this latter case it is easy to verify that the density function is continuous and that the derivatives match at the boundaries of each segment.

The distribution functions F_n do depend on the density ϕ in an important fashion. We interpret ϕ as a description of the individual's knowledge. When ϕ is diffuse over $[0,1]$ (e.g. uniform) all points on the simplex Δ_n are equally likely and in some instances the individual will have the belief that she can answer the question correctly while in some cases the beliefs will be relatively uninformative, while if ϕ is concentrated around, or in the limit at, ½ the individual is always indecisive. Finally, when ϕ's modes are located around 0 and 1, the individual is always relatively informed about the correct answer.

A. 2 Lifelines

Extending the model above to allow for the lifelines makes the analysis more difficult but also enables us to exploit more aspects of the data. We show how, in the first sub-section below, the model can be modified when only one lifeline is allowed for. In a second sub-section we show how the model can be modified for all three lifelines. We then present the precise assumptions that allow the modelling of each lifeline in particular.

A. 2.1 The complete game

We now assume that the three lifelines are available and that each can be played at most once. As above each lifeline generates a new belief **q** which is used in the

28

decision process instead of the individual's initial belief. Given the initial belief $\mathbf{p}$, the new belief is drawn from a separate distribution for each lifeline, say $H_1(\mathbf{q}|\mathbf{p})$ for "50:50", $H_2(\mathbf{q}|\mathbf{p})$ for "Ask The Audience" and $H_3(\mathbf{q}|\mathbf{p})$ for "Phone A Friend".

We write $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \gamma_3)$ for the "lifeline state" vector where $\gamma_i = 0$ if the i^{th} lifeline has been played and 1 otherwise. We use $W_n(\mathbf{p}; \boldsymbol{\gamma})$ to denote the optimal expected value of the game at stage n, when the probability vector of the current question is $\mathbf{p}$ and the lifeline state is $(\gamma_1, \gamma_2, \gamma_3)$. As above, $V_n(\mathbf{p})$ is used as a shorthand for $W_n(\mathbf{p}; (0,0,0))$ and $V_n(\mathbf{p})$ satisfies the recursive dynamic programming equations set out in section 5.1 above. Below we write the dynamic programming equations using the notation:

$$f_n((\gamma_1, \gamma_2, \gamma_3)) = \mathrm{E}[W_n(\mathbf{P}_n; (\gamma_1, \gamma_2, \gamma_3))], \tag{26}$$

where the expectation is with respect to $\mathbf{P}_n$, the distribution of the belief vector $\mathbf{p}$ at stage n.

When there are one or more lifeline left, i.e. $\gamma_1 + \gamma_2 + \gamma_3 \geq 1$, the contestant has three options: (i) quit, (ii) answer the question, (iii) use one of the remaining lifelines. The recursive equation is

$$W_n(\mathbf{p}; \boldsymbol{\gamma}) = \max\{a_{n-1}, p(f_{n-1}(\boldsymbol{\gamma}) - b_n) + b_n, k_n(\mathbf{p}; \boldsymbol{\gamma})\} \tag{27}$$

where $k_n(\mathbf{p}; \boldsymbol{\gamma})$ denotes the maximum expected value from using a lifeline when the belief is p and the lifeline state vector is $\boldsymbol{\gamma}$. Here,

$$k_n(\mathbf{p}; \boldsymbol{\gamma}) = \max_{i \in I(\boldsymbol{\gamma})} \left\{ \mathrm{E}\left[W_n(\mathbf{Q}_i; \boldsymbol{\gamma} - \mathbf{e}_i)|\mathbf{p}\right] \right\} \tag{28}$$

where $I(\boldsymbol{\gamma}) = \{j : \mathbf{e}_j \leq \boldsymbol{\gamma}\}$ using $\mathbf{e}_i$ to denote the i-th unit vector in R^3 and $\mathbf{Q}_i$ is distributed according to $H_i(\mathbf{q}|\mathbf{p})$. This formulation does not preclude an individual from using more than one lifeline on the same question, a behaviour we observe in some contestants.

A. 2.2 "50:50"

This is the simplest lifeline to model. It provides the contestant with "perfect information" since two incorrect answers are removed. Ex-ante (i.e. before the lifeline is played) the contestant believes that the correct answer is i (=1,…,4) with probability p_i. The "50:50" lifeline removes two of the incorrect answers, retaining $j \neq i$, say, with equal probability (1/3). By Bayes Theorem, the probability that answers i,j survive this elimination process is $p_i/3$. The answers i and j can also be retained if j is correct and i survives elimination. This occurs with probability $p_j/3$. Applying Bayes Theorem gives the updated belief vector $\mathbf{q}^{\{i,j\}}$, where

$$\mathbf{q}_k^{\{i,j\}} = \begin{cases} \dfrac{p_i}{p_i + p_j} & \text{if } k = i, \\[2mm] \dfrac{p_j}{p_i + p_j} & \text{if } k = j \\[2mm] 0 & \text{otherwise.} \end{cases} \tag{29}$$

Hence $H_1(\mathbf{q};\mathbf{p})$ is a discrete distribution with support $\left\{\mathbf{q}^{\{i,j\}}\right\}_{\{i,j\}\in\{1,2,3,4\}}$ and such that

$H_1\left(\mathbf{q}^{\{i,j\}};\mathbf{p}\right) = \left(p_i + p_j\right)/3$, and 0 elsewhere.

A. 2.3 "Ask the Audience"

Modeling the "Ask the Audience" lifeline requires more than simply applying Bayes' rule to the current belief draw. In particular we must allow the contestant to learn from the information provided by the lifeline, i.e. here the proportions of the audience's votes in favour of each alternative answer. The difficulty here is to understand why and how should a "perfectly informed" rational individual revise his/her prior on the basis of someone else's opinion? The route we follow here was proposed by French (1980) in the context of belief updating after the opinion of an expert is made available. French suggests that the updated belief that some event A is realised after some information inf has been revealed should be obtained from the initial belief, $\Pr[A]$, the marginal probability that a given realisation of the information is revealed, $\Pr[\text{inf}]$, and the

individual's belief about the likelihood that the information will arise if A subsequently occurs, $\Pr[\inf \mid A]$ according to the following rule, related to Bayes theorem:

$$\Pr[A \mid \inf] = \Pr[\inf \mid A]\Pr[A]/\Pr[\inf]. \tag{30}$$

In this expression $\Pr[\inf \mid A]$ is understood as another component of the individual's belief - her assessment of the likelihood of the signal given that the relevant event subsequently occurs.

Introducing $\overline{A}$, A's alternative event, this is rewritten as

$$\Pr[A \mid \inf] = \frac{\Pr[\inf \mid A]\Pr[A]}{\Pr[\inf \mid A]\Pr[A] + \Pr\left[\inf \mid \overline{A}\right]\Pr[\overline{A}]}. \tag{31}$$

In our context we understand the asking the audience as an appeal to an expert, and assume that the events of interest are the four events "answer k is correct", k=1,2,3,4. We assume that contestants "learn" some information about the quality of the expert in particular the distribution of the quantities

$$\Pr\left[\mathbf{q} = (q_1, q_2, q_3, q_4) \mid \text{answer } k \text{ is correct}\right] \equiv \theta_k, \tag{32}$$

where q_k is the proportion of votes allocated to the k^{th} alternative. Following French's proposal, the k^{th} component of the updated belief π given the information q is:

$$\pi_k = \theta_k p_k \Big/ \sum_{j=1}^{4} \theta_j p_j . \tag{33}$$

Let us assume for now that each contestant knows the joint distribution of the vector $\theta = (\theta_1, \theta_2, \theta_3, \theta_4)$. In fact the above expression implies that, without loss of generality, we can normalise the θ_k to sum to one. Denote $I(\theta)$ the density function of θ given some initial belief p. Given p, the density of the updated belief $H_2(\pi; \mathbf{p})$ can be calculated as:

$$H_2(\pi; \mathbf{p}) = I\big(\theta(\pi; \mathbf{p})\big) \left(\prod_{k=1}^{4} p_k\right) \left(\sum_{k=1}^{4} \pi_k p_k^{-1}\right)^4, \tag{34}$$

with $\theta_i(\pi;\mathbf{p}) = \pi_i p_i^{-1} \Big/ \sum_{k=1}^{4} \pi_k p_k^{-1}$. The term $\left(\prod_{k=1}^{4} p_k\right)\left(\sum_{k=1}^{4} \dfrac{\pi_k}{p_k}\right)^{4}$ arises because of the change of variable from θ to π.

The quantities $\Pr\big[\mathbf{q} = (q_1,q_2,q_3,q_4) \mid \text{answer } k \text{ is correct}\big] \equiv \theta_k$ represent the added information obtained from using the lifeline and are estimable from the data provided we assume a form of conditional independence. In particular we require that the contestant's choice to ask the audience does not influence the audience's answer. Furthermore, we assume that there is no information contained in the position of the correct answer, hence we expect the following symmetry restrictions to hold :

$$\Pr\big[\mathbf{q} = (q_1,q_2,q_3,q_4) \mid \text{answer 1 is correct}\big] = \Pr\big[\mathbf{q} = \big(q_{\sigma(1)},q_1,q_{\sigma(2)},q_{\sigma(3)}\big) \mid \text{answer 2 is correct}\big]$$
$$= \Pr\big[\mathbf{q} = \big(q_{\sigma'(1)},q_{\sigma'(2)},q_1,q_{\sigma'(3)}\big) \mid \text{answer 3 is correct}\big]$$
$$= \Pr\big[\mathbf{q} = \big(q_{\sigma''(1)},q_{\sigma''(2)},q_{\sigma''(3)}q_1,\big) \mid \text{answer 4 is correct}\big],$$

where $(\sigma(1),\sigma(2),\sigma(3))$, $(\sigma'(1),\sigma'(2),\sigma'(3))$ and $(\sigma''(1),\sigma''(2),\sigma''(3))$ are some permutations of $(2,3,4)$. The symmetry restrictions, the conditional independence assumption, and the uniform random allocation of the correct answer among four alternative answers allow us to estimate the likelihood of the information given the position of the correct answer, and therefore provide empirical estimates for $\Pr\big[\mathbf{q} = (q_1,q_2,q_3,q_4) \mid \text{answer } k \text{ is correct}\big]$.

In practice we assume that, given answer k is correct, information $\mathbf{q}$ has a Dirichlet density $D\big(\mathbf{q};\gamma_k(\lambda,\nu)\big)$, $k=1\ldots4$, defined over Δ_4 such that

$$D\big(\mathbf{q};\gamma_k(\lambda,\nu)\big) = \frac{\Gamma(3\nu+\lambda)}{\Gamma(\lambda)\Gamma(\nu)^3}\left(\prod_{i=1}^{4} q_i^{\nu-1}\right)q_k^{\lambda-\nu},$$

where the symmetry assumption is imposed through the parameter vector $\gamma_k(\lambda,\nu) = \nu + \mathbf{e}_k(\lambda-\nu)$ with $\mathbf{e}_k$ is a vector of zeros with a 1 in position k. This vector of parameters for the Dirichlet density depends on two free parameters only, λ and ν. These two parameters can be estimated (independently from the other parameters of the model) by maximum likelihood from the observation of the information obtained from

the audience (i.e. the histograms) whenever the lifeline is used, and the observation of which answer is the correct answer[12]. For completeness note that θ_k can be defined in terms of the elements of q as $\theta_k = q_k^{\lambda-\nu} \Big/ \sum_{j=1}^{4} q_j^{\lambda-\nu}$. The information density which the contestant expects is therefore the mixture $D(\mathbf{q};\mathbf{p},\lambda,\nu)$ of the previous densities $D(\mathbf{q};\gamma_k(\lambda,\nu))$, k=1…4, conditional on a given answer being correct, we have:

$$D(\mathbf{q};\mathbf{p},\lambda,\nu) = \sum_{i=1}^{4} p_i D(\mathbf{q};\gamma_i(\lambda,\nu)) = \frac{\Gamma(3\nu+\lambda)}{\Gamma(\lambda)\Gamma(\nu)^3}\left(\prod_{i=1}^{4} q_i^{\nu-1}\right)\left(\sum_{i=1}^{4} p_i q_i^{\lambda-\nu}\right), \tag{35}$$

where the mixing weights are the initial beliefs $p_i, i = 1...4$.

A. 2.4 "Phone a Friend"

To use this lifeline the contestant determines, ahead of the game, six potential experts ("friends") and when she plays the lifeline she chooses one from this list of six. We imagine that the contestant engages in some diversification when drawing up the list (i.e. the range and quality of "expert knowledge" of the friends on the list is in some way optimised), and at the time of playing the lifeline the contestant chooses the expert to call optimally.

There is however little information available to us about this process. As a consequence our model for this particular lifeline is somewhat crude. We assume that the entire process can be modelled as an appeal to an expert who knows the answer with some probability κ, and is ignorant with the probability $1-\kappa$. We assume that the expert informs the contestant of his confidence[13]. Hence either the contestant knows the answer and her opinion "swamps" the contestant's belief, or the expert is ignorant and conveys no information and the contestant's belief is left unchanged. The density of the updated belief is therefore:

$$H_3(\pi;\mathbf{p}) = \kappa \mathbf{1}_{[\pi=(1,0,0,0)]} + (1-\kappa)\mathbf{1}_{[\pi=\mathbf{p}]}. \tag{36}$$

---

[12] Even when the contestant chooses an incorrect answer, the compere always reveals the correct one.

[13] In practice, contestants invariably ask the friend how confident they feel – although the answer is usually not quantitative.

A. 3 Proposition (**factorisation of** $\chi_4(x_1, x_2, x_3, x_4)$)):

The joint density: $\chi_4(x_1, x_2, x_3, x_4) = \dfrac{2}{\mu_2} \phi(x_1) \phi\left(\dfrac{x_2}{1-x_1}\right) \phi\left(\dfrac{x_3}{1-x_1-x_2}\right)$, with

(x_1, x_2, x_3, x_4) such that $\displaystyle\sum_{i=1}^{4} x_i = 1$, $x_i \geq 0$ for all i, can be factorised as follows:

$$\chi_4(x_1, x_2, x_3, x_4) = f_{U_1}(x_1) \, f_{U_2|U_1}(x_2; x_1) \, f_{U_3|U_1,U_2}(x_3; x_1, x_2),$$

with $f_{U_1}(u)$, $f_{U_2|U_1}(v; u)$, $f_{U_3|U_1,U_2}(w; u, v)$, (conditional) densities such that

$$f_{U_1}(u) = \frac{(1-u)^2 \phi(u)}{\mu_2} \mathbf{1}_{[0 \leq u \leq 1]}, \text{ with } \mu_2 = \int_0^1 (1-x)^2 \phi(x)\, dx,$$

$$f_{U_2|U_1}(v; u) = 2 \frac{(1-u-v)}{(1-u)^2} \phi\left(\frac{v}{1-u}\right) \mathbf{1}_{[0 \leq v \leq 1-u]},$$

$$f_{U_3|U_1,U_2}(w; u, v) = \frac{1}{1-u-v} \phi\left(\frac{w}{1-u-v}\right) \mathbf{1}_{[0 \leq w \leq 1-u-v]}.$$

Proof: It is easy to verify, by simple integration for $f_{U_1}(u)$, $f_{U_2|U_1}(v; u)$, and by construction for $f_{U_3|U_1,U_2}(w; u, v)$, all three are well defined densities over the relevant ranges. Moreover their product is equal to $\chi_4(.)$.

This implies that if U_1, U_2 and U_3 are three random variables each distributed with densities $f_{U_1}(u)$, $f_{U_2|U_1}(v; u)$, and $f_{U_3|U_1,U_2}(w; u, v)$, then the random vector $P = \begin{pmatrix} U_1 & \bar{U}_1 U_2 & \bar{U}_1 \bar{U}_2 U_3 & \bar{U}_1 \bar{U}_2 \bar{U}_3 \end{pmatrix}$, with $\bar{U}_i = 1 - U_i$ for all i=1..3, is distributed with joint density: $\chi_4(x_1, x_2, x_3, x_4) = \dfrac{2}{\mu_2} \phi(x_1) \phi\left(\dfrac{x_2}{1-x_1}\right) \phi\left(\dfrac{x_3}{1-x_1-x_2}\right)$. Note that by construction $P'\mathbf{e} = 1$, and $P \geq 0$.

Since $\chi_4(\mathbf{x})$ and $\psi_4(\mathbf{x})$ share the same joint density for the order statistics, i.e. $4! \psi_4(\tilde{\mathbf{x}})$ where $\tilde{\mathbf{x}}$ is such that its element are sorted in descending order, to sample from $4! \psi_4(\tilde{\mathbf{x}})$ we propose to sample first from $\chi_4(.)$ and then to sort the resulting vector in descending order.

A.4 Probabilities and Simulated Likelihood

In this section we describe the evaluation of some of the probabilities that lead to the log likelihood. A complete description of the calculations can be obtained from the authors online[14].

Calculating the probabilities when only one lifeline is available.

When the contestant has used all her lifelines, the events of interest are the occurrences of the contestant quitting or losing, and for the last question the event that the contestant wins the million prize. The probabilities of these events can be calculated directly from the analytical expressions given in section 5.1 using the formulation for F we derive in section 5.3.

When one or more lifelines are available the calculations are made more complicated because of the information which is gained when the lifeline is used and which allows the contestant to update her belief. Hence, given the initial draw of the belief we determine whether this particular draw leads to the use of the lifeline and, if the lifeline is played, then whether the updated belief, or the original belief, if the lifeline is not played, is informative enough to lead the contestant to attempt an answer. Finally we evaluate the probability that the answer is correct (under the original or the updated belief).

We will write $\Omega_{k,n}^{ijk}(\mathbf{p})$ as the probability that given $\mathbf{p}$ at stage n event k (which is defined precisely below) is observed, given that the contestant is in the lifeline-state ijk, where i, (respectively j or k) is one if the first (respectively second or third) lifeline is yet to be played and zero otherwise. Let, $\Omega_{k,n}^{ijk}$ be the expected value of $\Omega_{k,n}^{ijk}(\mathbf{p})$ over all possible realisations of $\mathbf{p}$, i.e. $\Omega_{k,n}^{ijk} = \mathrm{E}\left[\Omega_{k,n}^{ijk}(\mathbf{P})\right]$. Finally $\Omega_{k,n}^{ijk,i'j'k'}(\mathbf{p})$ stand for the probability that given $\mathbf{p}$ at stage n event k is observed given that the contestant starts the question in the lifeline-state ijk and transit to lifeline-state $i'j'k'$. We consider below representative events for each lifeline.

---

[14] http://www.qub.ac.uk/schools/SchoolofManagementandEconomics/Staff/LanotGauthier/

"50:50" is the only lifeline available at stage n.

<u>The contestant uses "50:50", plays and wins </u>(moves to the next stage or wins the million prize).

First, define the probability that the contestant uses "50:50", plays and wins given a draw (ordered in decreasing order) **p** from the belief distribution:

$$
\Pr\left[\left\{\text{use "50:50"}\right\}\wedge\left\{\text{plays}\right\}\wedge\left\{\text{wins}\right\}\mid\left\{\text{stage n}\right\},\mathbf{p}\right]\equiv\Omega_{1,n}^{100}(\mathbf{p})=
$$
$$
\mathbf{1}\left[k_{n}^{1}(\mathbf{p},0,0,0)\geq p_{1}\left(f_{n-1}(1,0,0)-b_{n}\right)+b_{n}\right]\Omega_{1,n}^{100,000}(\mathbf{p}),
$$
(37)

where $\Omega_{1,n}^{100,000}(\mathbf{p})=\dfrac{1}{3}\displaystyle\sum_{j=1}^{3}p_{j}\sum_{k=j+1}^{4}\mathbf{1}\left[\pi_{jk}(\mathbf{p})\left(f_{n-1}(0,0,0)-b_{n}\right)+b_{n}\geq a_{n+1}\right]$, with $\pi_{jk}=\dfrac{p_{j}}{p_{j}+p_{k}}$.

This last expression is the probability that, given **p**, the contestant answers correctly after using the lifeline. Hence the unconditional probability satisfies

$$
\Pr\left[\left\{\text{use "50:50"}\right\}\wedge\left\{\text{plays}\right\}\wedge\left\{\text{wins}\right\}\mid\left\{\text{stage n}\right\}\right]\equiv\Omega_{1,n}^{100}=
$$
$$
\int_{\tilde{\Delta}_{4}}\mathbf{1}\left[k_{n}^{1}(\mathbf{p},0,0,0)\geq p_{1}\left(f_{n-1}(1,0,0)-b_{n}\right)+b_{n}\right]\Omega_{1,n}^{100,000}(\mathbf{p})\tilde{\psi}_{4}(\mathbf{p})d\mathbf{p},
$$
(38)

where $\tilde{\Delta}_{4}$ is the subset of the 4-simplex where $p_{1}\geq p_{2}\geq p_{3}\geq p_{4}\geq 0$.

In order to determine the probabilities we have used the fact that a contestant with a lifeline available will either use it (and perhaps then quit), or play. It is straightforward to verify that the five expressions above sum to unity; in particular the sum of the first three expressions is the probability that the contestant uses the lifeline and this is the complement of the sum of the last two probabilities.

Each term of the sum that determines $\Omega_{1}^{100}(\mathbf{p})$ (and similarly $\Omega_{2}^{100}(\mathbf{p})$ and $\Omega_{3}^{100}(\mathbf{p})$) is the product of the probability that a given two of the four options remain after the lifeline is played, with probability $\dfrac{1}{3}(p_{j}+p_{k})$, multiplied by the probability that the remaining alternative with the largest updated belief is correct, with probability $\pi_{jk}(\mathbf{p})=\dfrac{p_{j}}{p_{j}+p_{k}}$ with $p_{j}\geq p_{k}$, multiplied by the indicator that, given the updated belief, the contestant decides to play.

36

"Ask the Audience" is the only lifeline left at stage n.

The contestant uses "Ask the Audience", plays and loses,

$$\Pr\Big[\{\text{use "Ask the Audience"}\} \wedge \{\text{plays}\} \wedge \{\text{loses}\} \,|\, \{\text{stage } n\}\Big] \equiv \Omega_{2,n}^{010} =$$
$$\int_{\tilde{\Delta}_4} \Omega_{2,n}^{010}(\mathbf{p})\tilde{\psi}_4(\mathbf{p})\,d\mathbf{p},$$

(39)

where $\Omega_{2,n}^{010}(\mathbf{p}) = \mathbf{1}\Big[k_n^2(\mathbf{p},0,0,0) \geq p_1\big(f_{n-1}(0,1,0) - b_n\big) + b_n\Big]\Omega_{2,n}^{010,000}(\mathbf{p})$,

and $\Omega_{2,n}^{010,000}(\mathbf{p}) = \int_{\Delta_4}\big(1 - \pi_1(\mathbf{q};\mathbf{p})\big)\mathbf{1}\Big[\pi_1(\mathbf{q};\mathbf{p})\big(f_{n-1}(0,0,0) - b_n\big) + b_n \geq a_{n+1}\Big]D(\mathbf{q};\mathbf{p},\lambda,\nu)\,d\mathbf{q}$

where $\pi(\mathbf{q};\mathbf{p})$ stands for the revised belief after information vector $\mathbf{q}$ is made available and $\pi_1(\mathbf{q};\mathbf{p})$ is the largest element in $\pi(\mathbf{q};\mathbf{p})$.

"Phone a Friend" is the only lifeline left at stage n.

The contestant uses "Phone a Friend" and quits,

$$\Pr\Big[\{\text{uses "Phone a Friend"}\} \wedge \{\text{quits}\} \,|\, \{\text{stage } n\}\Big] \equiv \Omega_{3,n}^{001} = \int_{\tilde{\Delta}_4}\Omega_{3,n}^{001}(\mathbf{p})\tilde{\psi}_4(\mathbf{p})\,d\mathbf{p},$$

where $\Omega_{3,n}^{001}(\mathbf{p}) = \mathbf{1}\Big[k_n^3(\mathbf{p},0,0,0) \geq p_1\big(f_{n-1}(0,0,1) - b_n\big) + b_n\Big]\Omega_{3,n}^{001,000}(\mathbf{p})$

and $\Omega_{3,n}^{001,000}(\mathbf{p}) = (1 - \kappa)\mathbf{1}\Big[p_1\big(f_{n-1}(0,0,0) - b_n\big) + b_n < a_{n+1}\Big]$.

General Case: all the lifelines are available

When more than one lifeline is available at a given stage, the number of elementary events of interest increases, since not only can the contestants decide to play one lifeline among many but the contestant can play more than one lifeline to answer a single question. Hence while there are only five elementary events of interest when only one given lifeline is left there are nine such events when two lifelines are available and seventeen when all three lifelines are available, ignoring the order in which the contestant uses the lifeline and not counting events with zero probability ex-ante (for example observing an event such as quitting while the three lifelines are available)[15]. In this

---

[15] In the case of two lifelines left : 1) Uses the two lifelines, plays and wins; 2) Uses the two lifelines, plays and loses; 3) Uses the two lifelines, plays and loses; 4) Uses one of two lifelines, plays and wins; 5) Uses

section we present the relevant expressions needed to obtain the probabilities of few selected elementary event, all other probabilities can be obtained in a similar fashion.

The contestant uses the three lifelines (in any order), plays and loses.

$$\Pr\left[\{\text{uses all life lines}\}\wedge\{\text{plays}\}\wedge\{\text{loses}\}\,|\,\{\text{stage n}\}\right]\equiv\Omega_{2,n}^{111}=$$

$$\int_{\tilde{\Delta}_4}\begin{array}{l}\mathbf{1}\left[k_n^1(\mathbf{p},0,1,1)\geq\max\left\{p_1\left(f_{n-1}(1,1,1)-b_n\right)+b_n,k_n^2(\mathbf{p},1,0,1),k_n^3(\mathbf{p},1,1,0)\right\}\right]\Omega_{2,n}^{111,011}(\mathbf{p})\\[2mm]\mathbf{1}\left[k_n^2(\mathbf{p},1,0,1)\geq\max\left\{p_1\left(f_{n-1}(1,1,1)-b_n\right)+b_n,k_n^1(\mathbf{p},0,1,1),k_n^3(\mathbf{p},1,1,0)\right\}\right]\Omega_{2,n}^{111,101}(\mathbf{p})+\\[2mm]\mathbf{1}\left[k_n^3(\mathbf{p},1,1,0)\geq\max\left\{p_1\left(f_{n-1}(1,1,1)-b_n\right)+b_n,k_n^1(\mathbf{p},0,1,1),k_n^2(\mathbf{p},1,0,1)\right\}\right]\Omega_{2,n}^{111,110}(\mathbf{p})\end{array}\quad(40)$$

$$\tilde{\psi}_4(\mathbf{p})d\mathbf{p}.$$

where $\quad\Omega_{2,n}^{111,011}(\mathbf{p})=\dfrac{1}{3}\sum_{j=1}^{3}\sum_{k=j+1}^{4}\Omega_{2,n}^{011}\left(\left(\pi_{j,k}(\mathbf{p}),\pi_{k,j}(\mathbf{p}),0,0\right)\right),\qquad(41)$

$$\Omega_{2,n}^{111,101}(\mathbf{p})=\int_{\Delta_4}\Omega_{2,n}^{101}\left(\pi(\mathbf{q};\mathbf{p})\right)D(\mathbf{q};\mathbf{p},\lambda,\nu)d\mathbf{q},\qquad(42)$$

$$\Omega_{2,n}^{111,110}(\mathbf{p})=\kappa\Omega_{2,n}^{110}(1,0,0,0)+(1-\kappa)\Omega_{2,n}^{110}(\mathbf{p}).\qquad(43)$$

Inspection of these expressions reveals that the probabilities of events in which more than one lifeline is available, here $\Omega_{2,n}^{111}$, can be defined recursively in terms of the conditional probability of events with one fewer lifeline, given the initial belief draw, here $\Omega_{2,n}^{011}(\mathbf{p})$, $\Omega_{2,n}^{101}(\mathbf{p})$ and $\Omega_{2,n}^{110}(\mathbf{p})$. In turn, each of these conditional probabilities can be calculated from conditional probabilities involving only one lifeline, i.e. $\Omega_{2,n}^{001}(\mathbf{p})$, $\Omega_{2,n}^{100}(\mathbf{p})$ and $\Omega_{2,n}^{010}(\mathbf{p})$. This property is a consequence of the recursive definition of the value function over the lifeline part of the state space (see section 5.4.b).

Recall, however, that the number of events of interest when the three lifelines are available is larger than when only two or less are available. Hence the definition of 17

---

one of two lifelines, plays and loses; 6) Uses other lifeline, plays and wins; 7) Uses other lifeline, plays and loses; 8) Does not use any lifeline, play and win; 9) Does not use any lifeline, play and loses; ….
In the case of three lifeline left: 1) Uses the three lifelines, plays and wins; 2) Uses the three lifelines, plays and loses; 3) Uses the three lifelines, plays and loses; 4) Uses "50:50" and "Phone a Friend", plays and wins; 5) Uses "50:50" and "Phone a Friend", plays and loses; 6) Uses another "50:50" and "Ask the Audience", plays and win; 7) Uses "50:50" and "Ask the Audience", plays and loses;… ; 10) "Uses "50:50", plays and win; 11) Uses "50:50", plays and loses; …16) Does not use any lifeline, play and win; 17) Does not use any lifeline, play and loses;

probabilities with three lifeline at stage n, i.e. $\Omega_{m,n}^{111}$, $m=1\dots17$, will involve the 27 conditional probabilities with two lifelines, i.e. $\Omega_{m,n}^{011}(\mathbf{p})$, $\Omega_{m,n}^{101}(\mathbf{p})$ and $\Omega_{m,n}^{110}(\mathbf{p})$, $m=1\dots9$. In turn each of these conditional probabilities will depend on the 15 probabilities with one lifeline as defined in the previous section, i.e. $\Omega_{m,n}^{100}(\mathbf{p})$, $\Omega_{m,n}^{010}(\mathbf{p})$ and $\Omega_{m,n}^{001}(\mathbf{p})$ $m=1\dots5$.

The three lifelines are available, the contestant uses "50:50" , plays and wins.

$$\Pr\Big[\{\text{uses "50:50" only among 3 life lines}\}\wedge\{\text{plays}\}\wedge\{\text{wins}\}\,|\,\{\text{stage}\,n\}\Big]\equiv\Omega_{10,n}^{111}$$

$$=\int_{\tilde{\Delta}_4}\mathbf{1}\Big[k_n^1(\mathbf{p},0,1,1)\geq\max\Big\{p_1\big(f_{n-1}(1,1,1)-b_n\big)+b_n,k_n^2(\mathbf{p},1,0,1),k_n^3(\mathbf{p},1,1,0)\Big\}\Big]\Omega_{10,n}^{111,011}(\mathbf{p})\tilde{\psi}_4(\mathbf{p})d\mathbf{p}.$$

with $\Omega_{10,n}^{111,011}(\mathbf{p})=\dfrac{1}{3}\displaystyle\sum_{j=1}^{3}\sum_{k=j+1}^{4}\Omega_{8,n}^{011}\Big(\big(\pi_{j,k}(\mathbf{p}),\pi_{k,j}(\mathbf{p}),0,0\big)\Big)$ where $\Omega_{8,n}^{011}(\mathbf{p})$ is the probability that with "Ask the Audience" and "Phone a Friend" available, for some belief p, the individual plays and wins.

Three lifelines are available, the contestant does not use any, plays and loses.

$$\Pr\Big[\{\text{does not use any of the 3 life lines}\}\wedge\{\text{plays}\}\wedge\{\text{loses}\}\,|\,\{\text{stage}\,n\}\Big]\equiv\Omega_{17,n}^{111}$$

$$=\int_{\tilde{\Delta}_4}\mathbf{1}\Big[p_1\big(f_{n-1}(1,1,1)-b_n\big)+b_n\geq\max\Big\{k_n^1(\mathbf{p},0,1,1),k_n^2(\mathbf{p},1,0,1),k_n^3(\mathbf{p},1,1,0)\Big\}\Big](1-p_1)\tilde{\psi}_4(\mathbf{p})d\mathbf{p}.$$

Simulation and smoothing

The evaluation of the probabilities $\Omega_{m,n}^{rst}(\mathbf{p})$, $n=1..15$, $m=1..17$[16], $(r,s,t)\in\{0,1\}^3$ and of the conditional expectations $k_n^j(\mathbf{p},r,s,t)$, $n=1..15$, $j=1..3$, and $(r,s,t)\in\{0,1\}^3$ requires the use multidimensional integration techniques. Simulation methods (as described in Gouriéroux and Monfort (1996) and Train (2003)) are well suited and have been applied successfully in similar context (see the examples discussed in Adda and Cooper, (2003)).

Clearly the specification of the belief lends itself to a simulation based likelihood methodology since simulations of Beta variates are obtained simply from Gamma variates (see for example Poirier (1995)). In turn, Gamma variates can be obtained

---

[16] If $\Omega_{m,n}^{rst}$ is not defined for some m, and some r,s,t we assume $\Omega_{m,n}^{rst}=0$.

directly, using the inverse of the incomplete Gamma function. Numerically accurate methods to evaluate the inverse of the incomplete Gamma function are detailed in Didonato and Morris (1996)[17]. The main advantage of their results is that it allows for simulations that are continuous in the parameters of the Gamma distributions. Evaluation by simulation of an integral involving the density of a 4 dimensional Dirichlet random vector, $D(\mathbf{q};\mathbf{p},\lambda,\nu)$, is obtained directly by the simulation of each of its component. For example

$$\Omega_{1,n}^{100} = \int_{\tilde{\Delta}_4} \Omega_{1,n}^{100}(\mathbf{p})\tilde{\psi}_4(\mathbf{p})d\mathbf{p} = \int_{\tilde{\Delta}_4} \mathbf{1}_{\left[k_n^1(\mathbf{p},0,0,0)\geq p_1\left(f_{n-1}(1,0,0)-b_n\right)+b_n\right]}\Omega_{1,n}^{100,000}(\mathbf{p})\tilde{\psi}_4(\mathbf{p})d\mathbf{p}, \quad (44)$$

can be approximated by

$$\widehat{\Omega}_{1,n}^{100}(S) = \frac{1}{S}\sum_{s=1}^{S}\Omega_{1,n}^{100}(\mathbf{p}_s) = \frac{1}{S}\sum_{s=1}^{S}\mathbf{1}_{\left[k_n^1(\mathbf{p}_s,0,0,0)\geq p_{1,s}\left(f_{n-1}(1,0,0)-b_n\right)+b_n\right]}\Omega_{1,n}^{100,000}(\mathbf{p}_s), \quad (45)$$

where $\mathbf{p}_s$ is one of S (the number of simulations) independent draws from the distribution of the order statistics of the belief, $\tilde{\psi}_4(.)$. In fact the accuracy of this simulated probability (and of all others which involve draws from $\tilde{\psi}_4(.)$) can be improved upon through antithetic variance reduction techniques which involve the permutations of the gamma variates used to generate each individual beta variate[18] (as explained for example in Davidson and McKinnon (2004) or in Train (2003)). Moreover, the quantity

$$\Omega_{10,n}^{111} = \int_{\tilde{\Delta}_4} \mathbf{1}_{\left[k_n^1(\mathbf{p},0,1,1)\geq\max\left\{p_1\left(f_{n-1}(1,1,1)-b_n\right)+b_n,k_n^2(\mathbf{p},1,0,1),k_n^3(\mathbf{p},1,1,0)\right\}\right]}\Omega_{10,n}^{111,011}(\mathbf{p})\tilde{\psi}_4(\mathbf{p})d\mathbf{p}. \quad (46)$$

can be evaluated using the simpler formula

---

[17] This is implemented in Gauss in the procedure **gammaii** (contained in the file cdfchic.src).

[18] For example, to simulate a draw from a $B(\alpha,\alpha+2)$, one can draw two independent realisations of a random variable distributed according to a $\gamma(\alpha)$, say z_1 and z_2, and one realisation from a $\gamma(2)$, say z_3. Then both $z_1/(z_1+z_2+z_3)$ and $z_2/(z_1+z_2+z_3)$ are draws from a $B(\alpha,\alpha+2)$, furthermore they are negatively correlated, so that the variance of their mean is smaller than the variance of the mean of two uncorrelated draws from a $B(\alpha,\alpha+2)$. In fact the relative efficiency, measured by the ratio of the variances, is $\left(\frac{3\alpha+2}{4\alpha+2}\right)^2\left(\frac{3+3\alpha}{3+4\alpha}\right) < 1$ for $\alpha > 0$.

$$\hat{\Omega}_{10,n}^{111}(S) = \frac{1}{S}\sum_{s=1}^{S} \mathbf{1}_{\left[k_n^1(\mathbf{p}_s,0,1,1)\geq\max\left\{p_{1,s}\left(f_{n-1}(1,1,1)-b_n\right)+b_n,k_n^2(\mathbf{p}_s,1,0,1),k_n^3(\mathbf{p}_s,1,1,0)\right\}\right]}\Omega_{10,n}^{111,011}(\mathbf{p}_s), \qquad (47)$$

or any improvement of it. Similarly $\Omega_{2,n}^{111,101}(\mathbf{p}) = \int_{\Delta_4}\Omega_{2,n}^{101}\left(\pi(\mathbf{q};\mathbf{p})\right)D(\mathbf{q};\mathbf{p},\lambda,\nu)d\mathbf{q}$ can be

evaluated by $\hat{\Omega}_{2,n}^{111,101}(\mathbf{p};S) = \frac{1}{S}\sum_{i=1}^{4}p_i\sum_{s=1}^{S}\left[\Omega_{2,n}^{101}\left(\pi(\mathbf{q}_{s,i};\mathbf{p})\right)\right]$, where $\mathbf{q}_{s,i}$ is one of S independent

draws from $D\left(\mathbf{q};\gamma_i(\lambda,\nu)\right)$.

Finally all quantities $k_n^2(\mathbf{p},r,s,t) \equiv \mathrm{E}_{\pi_2|\mathbf{p}}\left[W_n\left(\underset{\sim}{\Pi},r,s,t\right)|\mathbf{p}\right]$ which involve a multi

dimensional integral and the joint density $D(\mathbf{q};\mathbf{p},\lambda,\nu)$ can be obtained in a similar

fashion: for example, using $\hat{k}_n^2(\mathbf{p},r,s,t;S) = \frac{1}{S}\sum_{i=1}^{4}p_i\sum_{s=1}^{S}W_n\left(\mathbf{q}_{s,i},r,s,t\right)$, where $\mathbf{q}_{s,i}$ is one of

S independent draws from $D\left(\mathbf{q};\gamma_i(\lambda,\nu)\right)$. In practice these expression are modified in

order to smooth out the discontinuities that are created by the indicator terms. Hence, the

indicator functions $\mathbf{1}\left[v_1 \geq \max\left\{v_2,v_3,v_4\right\}\right]$, $\mathbf{1}\left[v_1 \geq \max\left\{v_2,v_3\right\}\right]$, or $\mathbf{1}\left[v_1 \geq v_2\right]$, are replaced

by their smoothed versions, $\dfrac{1}{1+\exp\left(\eta(v_2-v_1)\right)+\exp\left(\eta(v_3-v_1)\right)+\exp\left(\eta(v_4-v_1)\right)}$,

$\dfrac{1}{1+\exp\left(\eta(v_2-v_1)\right)+\exp\left(\eta(v_3-v_1)\right)}$, and $\dfrac{1}{1+\exp\left(\eta(v_2-v_1)\right)}$ respectively, where η is a

smoothing constant. In the limit as $\eta \to +\infty$ the smoothed versions tend to the indicators.

Table 1   *Questionnaire Sample and Population Data*

	Population survey data*		WWTBAM competitors		WWTBAM non-competitors	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Male	0.52	0.40	0.76	0.43	0.66	0.48
Age	44.41	10.21	43.14	9.36	47.86	11.67
Married	0.80	0.44	0.79	0.41	0.76	0.43
Education years	13.88	4.10	13.71	3.99	12.82	3.22
Smoker [++]	0.25	0.42	0.22	0.41	0.26	0.44
Renter	0.25	0.33	0.144	0.35	0.177	0.38
Contents uninsured [+]	0.09	0.26	0.07	0.27	0.06	0.31
House value (£k) [**] [++]	178.9	157	190.8	127	184.8	188
Employed	0.652	0.44	0.638	0.48	0.593	0.49
Self-employed	0.155	0.38	0.193	0.40	0.189	0.39
Not working	0.194	0.40	0.160	0.37	0.195	0.40
Gross earnings (£k pa) [***]	27.08	23.0	31.17	24.0	28.67	22.7
Regular lottery player [+++]	0.67	0.40	0.63	0.41	0.65	0.41
Observations	various		243		548	

Notes: * the survey datasets have been re-weighted to reflect the gender mix in the WWTBAM data. Population data comes from the 2002 Labour Force Survey with the exception of: + from Family Expenditure Survey 2002 data, ++ from British Household Panel Study 2001 wave, and +++ from the Gambling Prevalence Survey 2002.   ** if owner occupier. *** if employed.

Table 2   *Questionnaire Contestant Sample and Population of Contestants*

	Questionnaire sample of contestants		Population of contestants on videotapes	
	Mean	Std Dev	Mean	Std Dev
Male	0.76	0.43	0.77	0.43
Winnings £,000	61.96	104.1	54.26	105.9
% quit last Q	0.68	0.47	0.67	0.47
N	243		515	

Note: We categorise players who won the maximum £1m as quitters.

Table 3   *Maximum Likelihood Estimates of the Parameters of the Distribution of Histograms (ATA)*

Parameter	Estimate	Std. err.
λ	4.754	0.210
ν	0.914	0.030
Number of observations	501	
Log-Likelihood	1526.41	

Note: These results are based on 501 observations for which the use of ATA is observed.

Table 4   Maximum Likelihood Estimates

$$u(x) = \frac{(x+\gamma)^{1-\rho}}{1-\rho} \qquad\qquad u(x) = \frac{-\text{sign}(\alpha)}{\alpha}\exp\left(-\frac{\alpha}{1-\rho}(x+\gamma)^{1-\rho}\right)$$

Parameters		Homogenous γ			Heterogenous γ (2 groups)		
	All	All	men	Women	all	Men	women
ρ	1.018	1.0567	1.0469	1.0778	1.0985	1.0783	1.1811
	(0.001)	(0.018)	(0.049)	(0.047)	(0.028)	(0.027)	(0.093)
γ_1	0.410	0.3138	0.3056	0.3879	0.0454	0.0425	0.0380
	(0.077)	(0.109)	(0.141)	(0.214)	(0.087)	(0.102)	(0.212)
γ_2	-	-	-	-	1.1238	0.9654	1.8072
					(0.321)	(0.356)	(0.899)
Prob[γ_1]	-	-	-	-	0.1416	0.1564	0.0958
α	-	-0.0309	-0.0377	-0.0278	-0.0343	-0.0345	-0.0411
		(0.006)	(0.019)	(0.006)	(0.003)	(0.005)	(0.019)
κ	0.419	0.354	0.364	0.320	0.529	0.534	0.484
	(0.027)	(0.024)	(0.029)	(0.045)	(0.024)	(0.028)	(0.048)
Mean log-lik	-9.28661	-9.28661	-9.35367	-9.03649	-8.82529	-8.89011	-8.57972
Number of Obs.	515	515	396	119	515	396	119

Note: Two further parameters are estimated. These parameters specify the dependence of the belief distribution on the question round.

Table 5   *Certainty Equivalents (£ ,000)*

		50:50		PAF		ATA	
a_{n+1}	b_n	CE	CE of LL	CE	CE of LL	CE	CE of LL
p=(0.9,0.05,0.03,0.02)							
500	32	701.85	848.12	701.85	912.84	701.85	845.07
250	32	409.16	435.58	604.16	484.06	522.76	438.82
125	32	232.61	225.51	335.19	241.31	282.85	229.18
64	32	139.39	121.47	187.43	120.31	157.20	123.14
32	32	89.86	73.34	114.35	74.13	91.43	73.95
16	1	47.12	46.08	54.37	53.63	46.91	47.35
8	1	26.08	19.72	35.06	22.49	26.44	20.22
4	1	14.04	11.05	19.70	12.15	13.88	11.30
2	1	8.95	7.22	12.70	7.69	9.30	7.35
1	1	6.46	5.38	8.78	5.64	6.78	5.46
0	0	4.35	4.04	5.62	4.29	4.47	4.15
p= (0.6,0.2,0.15,0.05)							
500	32	174.40	544.80	174.40	789.05	174.40	719.35
250	32	130.14	282.30	160.81	438.58	148.67	360.24
125	32	95.63	152.25	116.75	222.49	106.42	180.99
64	32	72.21	92.07	84.96	114.14	77.14	93.78
32	32	56.70	61.84	64.76	72.74	57.24	62.56
16	1	9.29	22.75	10.06	46.06	9.27	31.15
8	1	6.66	11.19	7.87	18.06	6.71	14.05
4	1	4.68	7.01	5.68	10.83	4.65	7.33
2	1	3.60	4.97	4.41	7.44	3.68	5.10
1	1	2.97	3.91	3.56	5.31	3.06	4.00
0	0	1.19	2.31	1.41	4.05	1.21	2.41
p= (0.4,0.3,0.2,0.1)							
500	32	123.81	499.10	123.81	780.24	123.81	719.35
250	32	98.08	254.91	116.07	396.12	109.05	360.24
125	32	76.73	133.95	89.95	217.27	83.55	181.05
64	32	61.33	81.53	69.82	110.96	64.65	93.03
32	32	50.56	57.46	56.22	72.13	50.94	59.99
16	1	6.16	18.08	6.57	46.92	6.14	31.15
8	1	4.70	9.22	5.38	18.65	4.73	14.05
4	1	3.52	5.73	4.12	10.81	3.50	7.37
2	1	2.84	4.22	3.35	7.31	2.89	4.64
1	1	2.43	3.40	2.81	5.29	2.49	3.70
0	0	0.81	1.80	0.94	3.85	0.82	2.09

Note: Certainty equivalent CE is calculated as $u^{-1}\left(p\left(f_{n-1}-b_n\right)+b_n\right)$, whereas the CE of LL is calculated as $u^{-1}\left(k_n\left(\mathbf{p},.,.,.\right)\right)$.

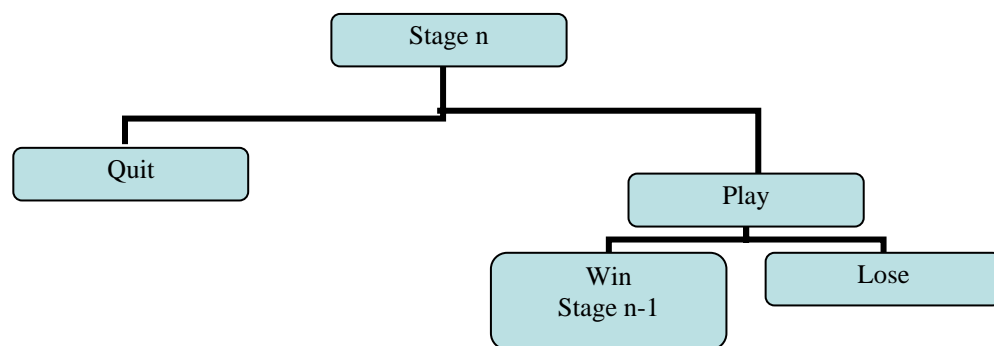*Figure 1a*   *Without Lifeline*

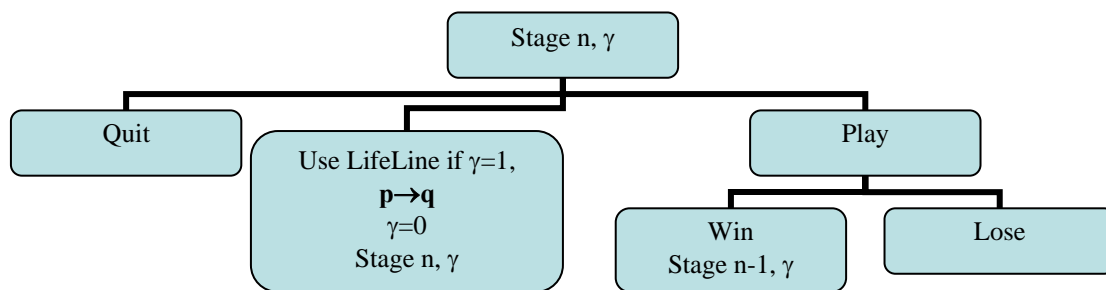*Figure 1b*   *With Lifeline*

45

Figure 2   *Distribution of winnings*

Figure 3   *Observed Fails and Quits Frequencies and Rates*

*Figure 4* *Distribution of the Maximum Belief*

Maximum Belief Distribution, F^4

Figure 5   *Value of playing the game at stage n, given the lifeline state.*
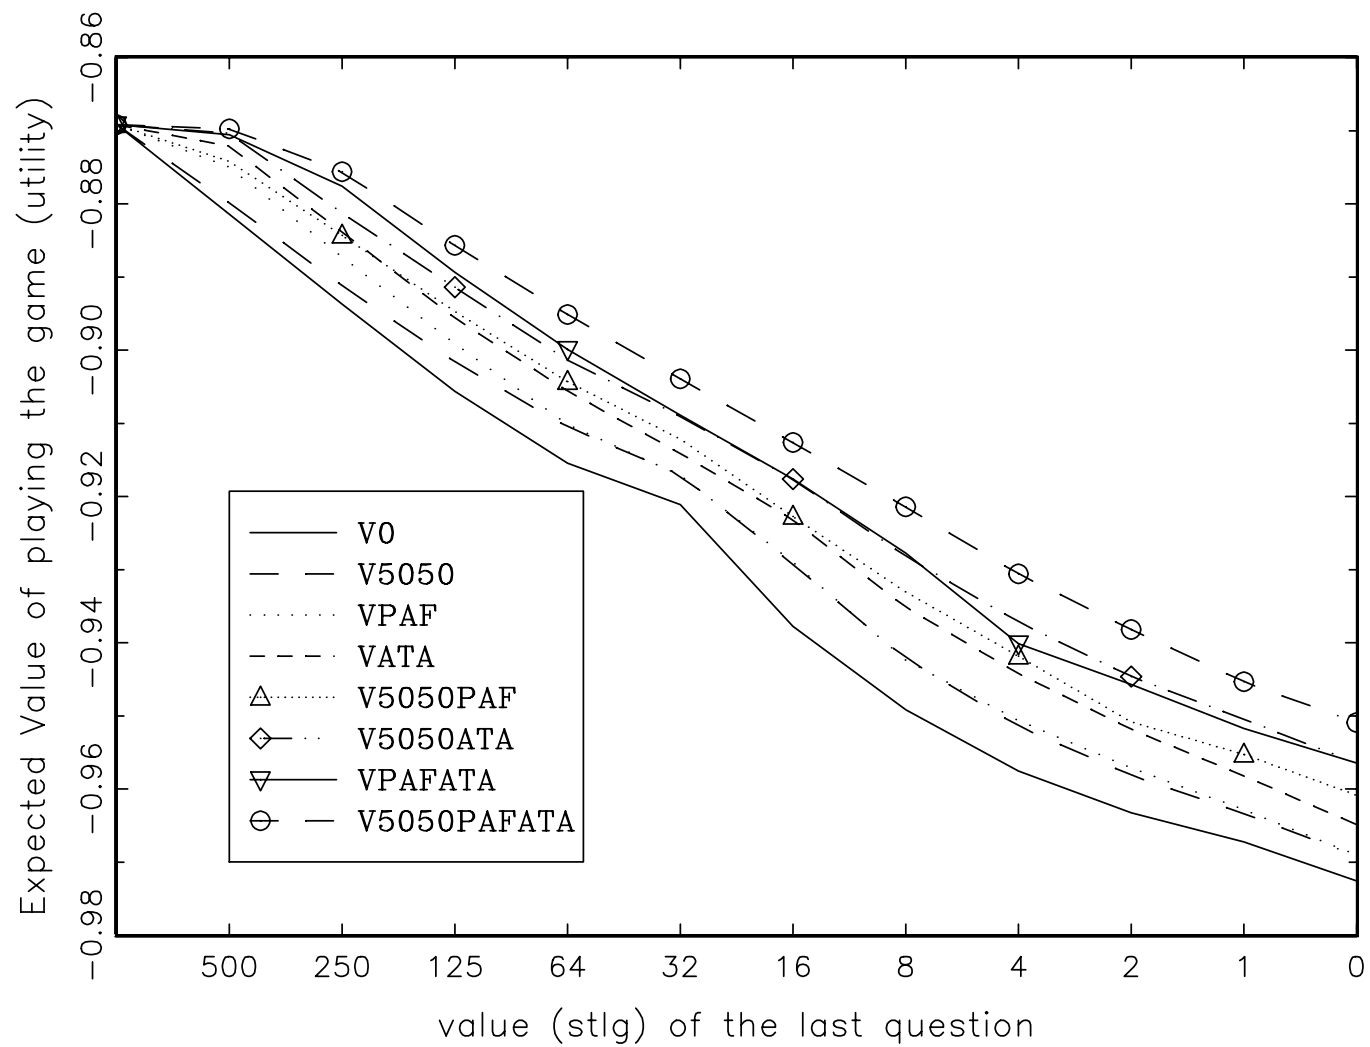
*Figure 6*   *Observed versus Predicted Frequencies of Fails and Quits*

50

*Figure 7:*   *Comparison of utility functions (All observations)*

Comparing the utility functions

Figure 8:   *Relative Risk Aversions (All observations)*