

Who Really Wants to be a Millionaire:
Estimates of Risk Aversion from Game Show Data

Roger Hartley

Gauthier Lanot

And

Ian Walker

No 719

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

Gauthier Lanot, Queen's University, Belfast, BT7 1NN

and

Ian Walker, University of Warwick, Coventry CV4 7AL

31 January 2005

Keywords: Risk aversion, gameshow

JEL Class: D810, C930, C230

Abstract

There is a considerable variation in estimates of the degree of risk aversion in the literature. This paper analyses the behaviour of contestants in one of the most popular TV gameshows ever to estimate a CRRA model of behaviour. 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 – up to £1 million. Our data sources have the virtue that we are able to check the representativeness of the gameshow participants. While the game requires skill, which complicates our analysis, the structure of the game is very simple so that complex probability calculations are not required of participants.

The CRRA model is complex despite its restrictiveness because of the sequential nature of this game – answering a question correctly opens the option to hear the next question and this has a value that depends on the stage of the game and the player's view about the difficulty of subsequent questions.

We use the data to estimate the degree of risk aversion and how it varies across individuals. We investigate a number of departures from this simple model including allowing the RRA parameter to vary by gender and age. Even though the model is extremely restrictive, in particular, it features a single RRA parameter we find that it fits the data across a wide range of wealth remarkably well and yields very plausible parameter values.

* We are grateful to Celador PLC, the gameshow's creators and the UK production company, for their help with gathering the data, and their permission to use it. We are particularly grateful to Ruth Settle who provided detailed advice that helped our understanding of the game. The research was supported by grant R000239740 from the Economic and Social Research Council. We are also indebted to seminar participants at the Institute for Fiscal Studies, Her Majesty's Treasury, the Welsh Assembly, CAM Copenhagen, the Universities of Bristol, Melbourne and Durham, and the Cardiff Business School. However, the views expressed herein are those of the authors alone.

1. Introduction

The idea that utility is concave and increasing in wealth is both a fundamental presumption in economics and important for a range of applied economic issues. For example, the extent of concavity, measured, say, by the degree of relative risk aversion has implications for any decisions that concern uncertainty – most obviously decisions such as insurance and portfolio choice, but less obviously decisions over phenomena such as self-employment, R&D, education, crime, smoking, and substance abuse¹.

This paper provides estimates of the degree of risk aversion using gameshow data, and addresses the main weaknesses of the existing literature. In particular, we provide estimates obtained from circumstances where the wealth of the same individuals varies dramatically and where individuals face very simple gains and losses. Thus, our evidence is well suited to addressing the question of whether expected utility theory can be applied across a wide range of wealth and whether constant relative risk aversion can be used to characterise behaviour in these circumstances.

The existing empirical literature that addresses the degree of risk aversion is distinguished by the breadth of its estimates and the literature is particularly sparse on attitudes towards large gambles. This is particularly troubling because an important deduction has been made by Rabin (2000) and Rabin and Thaler (1999) which could not be easily refuted with the available evidence. In particular they argue that expected utility maximising individuals are risk averse over small gambles and that any reasonable degree of risk aversion over small gambles would imply absurd choices over large gambles. Thus, they argue that it is not possible to characterise behaviour over a wide range of gambles using a CRRA Expected Utility function. Indeed, the authors use the argument to cast doubt on the very idea of expected utility maximisation.

This paper analyses the information on contestants in one of the most popular TV gameshows ever to estimate a CRRA model of behaviour. This gameshow has a number of features that makes it well suited for our analysis: the format is extremely straightforward and it involves no strategic decision-making, we have a large number

¹ Moreover, the cardinalisation of utility is important for public policy relating to optimal redistributive taxation (see Atkinson (1977)).

of observations, and the prizes are in cash paid immediately and cover a large range – up to £1 million. We use the data to estimate the degree of risk aversion and show how it varies across individuals. In particular, we investigate the extent to which the RRA parameter varies by gender and age. We also test the extent to which RRA may not be constant across wealth. Even though the model is extremely restrictive, in particular it features a constant RRA parameter across a wider range of wealth, we find that it fits the data remarkably well and yields very plausible parameter values.

Our data comes from the world's most popular TV gameshow of all time, *Who wants to be a millionaire?* (hereafter WWTBAM), which is played and broadcast under license in many countries but originated in the UK. 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 player is faced with a sequence of 15 multiple-choice questions. At each stage she can guess the answer to the current question and stands to double her current winnings but at the risk of losing a stage-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. Although there is no strategic element, contestants simply play against the house, it is a game where skill matters, which complicates our analysis, the structure of the game is very simple so that complex probability calculations are not required of participants.

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 very 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 data was transcribed from the original videotapes of the population of contestants. We further 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 could compare with population surveys such as the Labour Force Surveys.

We focus on CRRA preferences, despite its restrictiveness, because the sequential nature of the game gives rise to an 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 player’s view about the difficulty of subsequent questions, and the degree of risk aversion. This option value characteristic would complicate the econometric analysis considerably and the assumption of CRRA allows us to construct a model that can be estimated.

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 section 4 we provide a simple model of the game² that captures its formal structure so we can show the mechanics of the game in a straightforward way. We go on, in Section 5, to generalise this to embrace all the practical details of the game. In section 6 we present the econometric details and the likelihood.. In section 7 we give some summary details of the UK data and explain how we estimate risk aversion using this data. In section 8 we present some results and consider possible shortcomings of the work. In section 9 we concludes and draw together some conclusions, and outline some extensions of the work for the future.

2. Existing evidence

There are several distinct strands to the empirical literature. Firstly, because the coefficient of risk aversion enters into decisions that do not explicitly involve uncertainty but require cardinal utility, 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 constant relative risk aversion (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

² The game originates in the UK and the main difference across countries is in the units for the prizes and in their tax treatment. We hope to exploit the differences in prizes across countries, and across time within some countries, in future work.

is small. Thus, the macro-econometric literature largely suggests that the degree of risk aversion is large. Some of this literature³ considers two assets and backs out risk aversion 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 50⁴.

However, this method, that relies on portfolio allocations, has only ever been applied to microdata in a handful of studies. Attanasio, Banks and Tanner (2002) provides a very plausible estimate of the coefficient of relative risk aversion of just 1.44 using a large UK sample survey (for the sub-sample at an interior solution (i.e. of shareholders)), and appears to be unique in failing to reject the overidentifying restrictions implied by economic theory.

Jianakopulos 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 aversion 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 using tax registers for more than 7,000 households for. This study also recognizes the existence of real as well as financial assets and accounts 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 Jianakopulos and Bernasek but, in this case, not systematically correlated with characteristics apart 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 the coefficient of CRRA, in the atemporally separable case, depends on the ratio of income and wage

³ Notable contributions to this area are Epstein and Zin (1989, 1991).

⁴ 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.

elasticities and that the estimates in the labour supply literature implies a CRRA coefficient of about 1 and that a positive uncompensated wage elasticity is sufficient to bound CRRA 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 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 (a monthly premium of \$0.45 on average) relative to the expected loss (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 an average 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% change 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 actually correlated with risky behaviour in the data such as smoking, insurance and home ownership.

Donkers *et al* (2001) is a 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 (1980)). 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.

A further 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 CRRA are extremely large (more than 20) in each which might suggest that the questionnaire method is contaminating true risk aversion with some 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). However, there remains a worry that responses to hypothetical gambles are contaminated and do not reflect risk attitudes alone.

The present paper belongs firmly to the final strand to 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 depended on the expected value of the gamble from which he could deduce the degree of risk aversion. Given the rather small stakes involved, he found that the implied preferences⁵ were not significantly different from risk neutrality.

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 the coefficient of relative risk aversion, and the data again suggests that contestants are near risk neutral. Gertner (1993) analyses contestants in *Card Sharks* who enter a bonus round which involves a sequence of bets where the stakes are drawn from winnings in an earlier round that depends on the relative skill of contestants. He uses data on just the final bet in the bonus round and finds evidence of a high degree of risk aversion

⁵ They also model the ability of players 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.

(perhaps as high as 15), although he also found evidence of behaviour that contradicts expected utility theory.

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. One difficulty with this game is that prizes are annuities and so their value to players will depend on time preference. They find, assuming a discount rate of 10%, that contestants display risk-aversion with the mean ρ range from 0.64 to 1.76.

Finally, and closest to this study, Beetsma and Schotman (2001) use a Dutch game called *Lingo*. Like WWTBAM this is a game of skill. They use data from a television game show involving elementary lotteries as if it were a natural experiment so as to measure risk attitudes. Their dataset is large but the monetary stakes are, on average, relatively small. CRRA and CARA utility specifications are found to perform approximately equally well and they find robust evidence of a substantial degree of risk aversion with estimates of ρ in the range from 3 to 7. Extensions of the basic model, which allow for a separate utility flow purely from playing the game or for decisions based on decision weights instead of actual probabilities, raise the estimated degree of risk aversion.

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. In many of these countries the show was originally the most popular show on TV for some time. The game features a sequence of fifteen “multiple-choice” questions with associated prizes that, in the UK, start at £100 and (approximately) doubles each 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 players 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 difficulty of questions rises across the sequence of

questions⁶ and the fallback level 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, players may use one or more lifelines to improve their odds.

Contestants are not selected randomly 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 get asked a medium difficulty question. If correct their names get entered into a draw to appear in the studio. Ten names are drawn for each show. 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 to ensure that players 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.

4. A simple version of WWTBAM and a bound on risk aversion

It is useful to begin by considering a very simple model where utility exhibits constant relative risk aversion, there is only one question (think of this, for the moment, as being the last question in the sequence of questions faced by a contestant) and no lifelines (imagine they have all been used on earlier questions), and where the fallback level of winnings is fixed at some value, b . The purpose of this simple model is to introduce the game, highlight the issues, and use it to provide a crude idea of what the degree of risk aversion might be, before we attempt to construct, and estimate formally, a model that captures all of the complexities of the actual game. This stylised game can be characterised by

⁶ What “difficulty” means here is, of course, subjective. Many early questions are concerned with popular culture and sport. The details differ slightly from country to country but for the UK data used here the production staff divide questions into bins of what they regard as rising difficulty and contestants face a sequence of questions drawn randomly from successive bins. In early shows there were fewer than 15 bins because the production team did not have the experience to rank questions precisely. However, in recent shows there have been 15 such bins, one for each prize level.

$$\begin{aligned}
U(\textit{Quit}) &= \frac{1}{1-\rho} W^{1-\rho} \\
U(\textit{Gamble}) &= \frac{1}{1-\rho} [P 2^{1-\rho} W^{1-\rho} + (1-P)b^{1-\rho}]
\end{aligned}
\tag{1}$$

where W is the level of winnings in previous questions and P is the subjective probability of choosing the correct answer. We assume that the questions (i.e. the question itself as well as all candidate answers) are drawn randomly from a pool, such that, for any question, individuals are able to assign to each answer some subjective probability of being the correct answer. Hence P is a random variable with a distribution which is known to the individual. The decision problem of the contestant can be couched as a simple stopping rule: choose to answer the question if the subjective probability of being correct, P exceeds some critical value \bar{p} given by

$$\bar{p} = \frac{W^{1-\rho} - b^{1-\rho}}{2^{1-\rho} W^{1-\rho} - b^{1-\rho}}
\tag{2}$$

and otherwise quit and take the accumulated winnings W . The comparative statics of this simplified model suggests that individuals are more likely to quit the higher is ρ , and the lower (higher) is b according to $\rho > (<) 1$.

In the UK game, b is zero for the first 5 questions, £1000 for the next 5, and £32000 for the last five. In practice there are very few instances of quits or failures below £1000. Two probabilities are of particular interest: the probability of failure when the last question asked is worth £2,000 and the probability of failure when the last question asked is £64,000. In both cases these questions should be attempted by all individuals reaching that stage of the game because there is no downside risk at those stages of the game. Thus, straightforward estimates of the probability of a correct answer provide a measure the average difficulty of the questions. We focus here on the probability of failing at the £64,000 question which, in our data, is about $1/3^{\text{rd}}$.

To simplify matters further we assume that the £125,000 prize is 128,000 ($=2^7$) and we assume that the probability of failing this question can be approximated by the probability of failing the £64,000 question. In keeping with the simplified model outlined above a contestant will tackle this question whenever the expected utility of the gamble is greater than the utility of quitting with £64,000.

Two cases arise: $\rho > 1$ and $0 < \rho < 1$. If $\rho > 1$ the condition that determines the decision to answer this question becomes $\frac{2}{3} \cdot \frac{128^{1-\rho}}{(1-\rho)} + \frac{1}{3} \cdot \frac{32^{1-\rho}}{(1-\rho)} \geq \frac{64^{1-\rho}}{(1-\rho)}$. Since $1-\rho < 0$ this is equivalent to $\frac{2}{3} \cdot 128^{1-\rho} + \frac{1}{3} \cdot 32^{1-\rho} \leq 64^{1-\rho}$. Furthermore, since $128 = 2^7$, $64 = 2^6$, and $32 = 2^5$ this implies $\frac{1}{3} (2^{2(1-\rho)+1} + 1) \leq 2^{(1-\rho)}$. Substituting $2^{1-\rho} = \theta$, with θ varying from 1 to 0 as ρ varies from 1 to ∞ , we require $\frac{1}{3} (2\theta^2 + 1) \leq \theta$. Note that the equation $\frac{1}{3} (2\theta^2 + 1) - \theta = 0$ has one unique solution, θ^* , in the interval $[0,1]$. To the right of this solution the inequality above is satisfied, to the left it is not. Obviously, $\theta^* = 1/2$. Hence ρ must be such that $2^{(1-\rho)} \geq 1/2$, which gives $1 < \rho \leq 2$. That is, if $\rho > 1$ it must be the case that contestants cannot be too risk averse.

In the second case, $0 < \rho < 1$, and from the calculations above we find that ρ must be such that $\frac{1}{3} (2^{2(1-\rho)+1} + 1) \geq 2^{(1-\rho)}$. Substituting $2^{1-\rho} = \theta$, with θ varying from 2 to 1 as ρ varies from 0 to 1 we get $\frac{1}{3} (2\theta^2 + 1) \geq \theta$, which is satisfied for $\theta > 1$ and in turn implies $\rho < 1$ suggesting even more modest risk aversion.

On the basis of these bounds, and bearing in mind that the simplified model here ignores the option values of continuing, we conclude that it must be the case that $\rho < 2$ for play to continue beyond this level. The fact that many individuals are observed to play beyond the £64,000 question suggests that risk aversion is, in fact, quite low.

5. Extensions to the simplified version of WWTBAM

5.1. Dynamics

The model of participation we present now accounts for the potential future stages of the game, we focus on a simplified version of the game in which players are risk neutral and hence are expected income maximisers. We also ignore the “lifelines” of “asking the audience”, “phoning a friend” or “50:50” (which randomly discards 2

of the 3 incorrect answers) and assume that questions are selected by independent random drawing from a pool of questions of identical difficulty.

Let p denote the probability that the player (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 sections, a model for this distribution). Rounds of the game are denoted by the number of questions remaining, i.e. $n=N, \dots, 1$. Let a_n be the accumulated winnings after the player has successfully completed $N-n$ 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 winnings that are “protected”, i.e. the winnings that can be kept in the event of an incorrect answer. In the televised game the sequence of protected 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 player is faced with the first of 15 questions. The value of playing the game, and therefore answering the first question, is given by $V_{15}(p) = \max\{a_{16}, p(f_{14} - b_{14}) + b_{15}\}$ where $f_{14} = E[V_{14}(P)]$ is the optimal expected value of the next questions 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 with probability p , the value of the game is

$$V_n(p) = \max\{a_{n+1}, p(f_{n-1} - b_n) + b_n\}, \quad (3)$$

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

At any round of the game, there exists a critical value $\bar{p}_n = (a_{n+1} - b_n) / (f_{n-1} - b_n)$ such that if $p \leq \bar{p}_n$ the individual abandons 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 $V_n(p) = p(f_{n-1} - b_n) + b_n$. Hence the immediate

value of answering correctly is a_{n-1} and the expected difference $p(f_{n-1} - a_{n-1})$ represents the “option value” of continuing. These dynamic programming equations lead to the following relationship for the $\{f_n\}$:

$$f_{n-1} - f_n = (f_{n-1} - b_n) \int_{\bar{p}_n}^1 F(p) dp. \quad (4)$$

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 questions, is

$$\Pr[\text{"Play"}] = 1 - F(\bar{p}_n) \equiv \bar{F}(\bar{p}_n). \quad (5)$$

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

$$\Pr[\text{"Win"} | \text{"Play"}] = \frac{\int_{\bar{p}_n}^1 p dF(p)}{1 - F(\bar{p}_n)} \equiv \frac{\bar{G}(\bar{p}_n)}{\bar{F}(\bar{p}_n)}. \quad (6)$$

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

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

$$L(k, 0) = \{1 - \bar{F}(\bar{p}_k)\} \prod_{n=k+1}^{15} \bar{G}(\bar{p}_n). \quad (7)$$

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

$$L(k, 1) = \{\bar{F}(\bar{p}_k) - \bar{G}(\bar{p}_k)\} \prod_{n=k+1}^{15} \bar{G}(\bar{p}_n). \quad (8)$$

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

$$L(1, \cdot) = \prod_{n=1}^{15} \bar{G}(\bar{p}_n). \quad (9)$$

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 $u(x)$, consider $\{\tilde{a}_n\}_{i=1}^{16} = \{u(a_n)\}_{i=1}^{16}$ and $\{\tilde{b}_n\}_{i=1}^{15} = \{u(b_n)\}_{i=1}^{15}$ instead of $\{a_n\}_{i=1}^{16}$ and $\{b_n\}_{i=1}^{15}$.

5.2. Questions, Answers and Beliefs

The purpose of this section is to propose a model for the distribution of the beliefs that an individual hold each time a question and several answers (in the real game, four) are presented to her. In this section and in the next, we take as a given that the player 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 question/answer setting process we have in mind can be described as follows: first a given question and its possible answers in some specific order are drawn uniformly (at each stage of the game) from a pool of questions and corresponding candidate answers. The question and its possible answers (possibly in a different order) are presented to the candidate who is then endowed with a draw from the belief distribution concerning the likelihood of each answer. The formation of beliefs for all candidates is assumed to follow this process in an identical and independent manner. Hence, given a particular question, two identical individuals can hold distinct beliefs concerning the likelihood of each answer. Furthermore, any given individual can evaluate the distribution of her possible beliefs over the population of questions involved at any given stage of the game.

Formally, suppose that \mathbf{X} is an n -dimensional random vector with a continuous distribution on the simplex

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

We assume that \mathbf{X} has the probability density function $\psi_n(\mathbf{x})$, and we require it to exhibit the following symmetry property :

Let $\mathbf{x}, \tilde{\mathbf{x}} \in \Delta_n$, such that $\tilde{\mathbf{x}}$ is obtained from \mathbf{x} by any permutation of two distinct elements, then $\psi_n(\mathbf{x}) = \psi_n(\tilde{\mathbf{x}})$.

For our purposes we limit our investigation to the cases where $n \leq 4$. Our construction starts by considering a symmetric probability density function ϕ on $[0,1]$, i.e. such that $\phi(x_1) = \phi(1-x_1)$ for all x_1 in $[0,1]$. Note that ϕ 's symmetry implies

$$\int_0^1 \phi(x)(1-x) dx = \frac{1}{2} \text{ and } \int_0^1 \phi(x)(1-x)^2 dx = \int_0^1 \phi(1-x)(x)^2 dx = \mu_2 \quad (10)$$

where μ_2 is the second moment of ϕ . Note Φ the distribution function corresponding to ϕ . It is then straightforward to show that $\Phi[1-z] = 1 - \Phi[z]$.

Our construction of a class of belief distribution is based on ϕ . In the three cases of interest, we propose the following

$$\psi_2(x_1, x_2) = \frac{1}{2} [\phi(x_1) + \phi(x_2)], \quad (11)$$

$$\psi_3(x_1, x_2, x_3) = \frac{1}{3} \sum_{\{i,j,k\} \in \mathcal{P}_3} \phi(x_k) \phi\left(\frac{x_j}{1-x_k}\right), \quad (12)$$

$$\psi_4(x_1, x_2, x_3, x_4) = \frac{1}{12\mu_2} \sum_{\{i,j,k,l\} \in \mathcal{P}_4} \phi(x_l) \phi\left(\frac{x_k}{1-x_l}\right) \phi\left(\frac{x_j}{1-x_l-x_k}\right), \quad (13)$$

where, for any n , \mathcal{P}_n is the set of all permutations of $\{1, \dots, n\}$, 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 candidate answers are presented to the participant. Because ϕ is itself symmetric some (more or less obvious) simplifications are possible, we have

$$\psi_2(x_1, x_2) = \phi(x_1), \quad (14)$$

$$\psi_3(x_1, x_2, x_3) = \frac{2}{3} \left[\phi(x_2) \phi\left(\frac{x_2}{1-x_1}\right) + \phi(x_2) \phi\left(\frac{x_3}{1-x_2}\right) + \phi(x_3) \phi\left(\frac{x_1}{1-x_3}\right) \right], \quad (15)$$

$$\psi_4(x_1, x_2, x_3, x_4) = \frac{1}{6\mu_2} \sum_{\substack{l,k=1,\dots,4 \\ k \neq l}} \phi(x_l) \phi\left(\frac{x_k}{1-x_l}\right) \phi\left(\frac{x_j}{1-x_l-x_k}\right), \quad (16)$$

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 ordered statistics and for the distributions of the maximum amongst (x_1, \dots, x_n) .

5.3. Distribution of the maximum belief

The dynamic model outlined above involves the distribution, F , of the individual assessment on her chances of answering the question successfully. In the case without life lines, $F \equiv F_n$ 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. Hence, in this section we describe formulae for the distributions

$$F_n(z) \equiv \Pr\left[\bigcap_{i=1}^n \{X_i < z\}\right], \quad (17)$$

given that \mathbf{X} is distributed with density function ψ_n . In particular we can show (see appendix A for the details) that :

$$F_2(z) = \begin{cases} 0 & \text{if } z \leq \frac{1}{2} \\ 2\Phi[z] - 1 & \text{if } \frac{1}{2} \leq z \leq 1, \\ 1 & \text{otherwise.} \end{cases} \quad (18)$$

and as a consequence the density function $f_2(z)$, $z \in [0,1]$, is such that

$$f_2(z) = \begin{cases} 0 & \text{if } z \leq \frac{1}{2} \\ 2\phi[z] & \text{if } \frac{1}{2} \leq z \leq 1 \end{cases} \quad (19)$$

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

$$F_3(z) = 2 \int_{1-z}^1 F_2\left(\frac{z}{y}\right) y \phi(y) dy, \quad (20)$$

$$F_4(z) = \frac{1}{\mu_2} \int_{1-z}^1 F_3\left(\frac{z}{y}\right) y^2 \phi(y) dy, \quad (21)$$

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

$$F_4(z) = \begin{cases} 0 & \text{if } 0 \leq z \leq \frac{1}{4} \\ (4z-1)^3 & \text{if } \frac{1}{4} \leq z \leq \frac{1}{3} \\ -44z^3 + 60z^2 - 24z + 3 & \text{if } \frac{1}{3} \leq z \leq \frac{1}{2} \\ 1 - 4(1-z)^3 & \text{if } \frac{1}{2} \leq z \leq 1 \end{cases} \quad (22)$$

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.

Although the formula above tends to hide it, 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 $[0,1]$ (i.e. 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, $\frac{1}{2}$ 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.

5.4. Lifelines

Extending the model above to allow for the life lines makes the analysis more difficult but also enables us to exploit more aspects of the data. We first show how the model can be modified when only one life line is allowed for, and in a second subsection we show how the model is modified when all three life lines are included. We then present the precise assumptions that allow the modelling of each life line in particular.

5.4.a. A simplified game

Let us first consider the game with only one life line (say "50:50", although the discussion does not depend on the properties of "50:50"). To clarify the difference

figure 1 presents the decision trees at stage n and to stage with or without the life lines.

Hence, to account for the life line, the state space has to be extended. We will write $W_n(\mathbf{p};\gamma)$ 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 life line left and we will write the value function as $V_n(p;0)$.

In what follows we assume that the life line is a draw of a new belief, say \mathbf{q} , given \mathbf{p} the current belief. For example the use of the "50:50" reduces two components of the belief 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. Finally we define

$$k_n(\mathbf{p}) \equiv E_{\mathbf{q}|\mathbf{p}}[V_n(q,0)|\mathbf{p}], \quad (23)$$

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 following dynamic programming equations. When no lifeline is left we have the familiar equation:

$$V_n(p,0) = \max\{a_{n+1}, p(f_{n-1}(0) - b_n) + b_n\}, \quad (24)$$

where $f_n(0) = E[V_n(P;0)]$. With the life line left the contestant will choose the largest of the three options in the first choice line in Figure 1b, where :

$$W_n(\mathbf{p},1) = \max\{a_{n+1}, k_n(\mathbf{p}), p(f_{n-1}(1) - b_n) + b_n\}, \quad (25)$$

$f_n(1) \equiv E[W_n(\mathbf{P},1)]$ and $f_0(1) = f_0(0) = a_1$.

Figure 1a Without Life Line

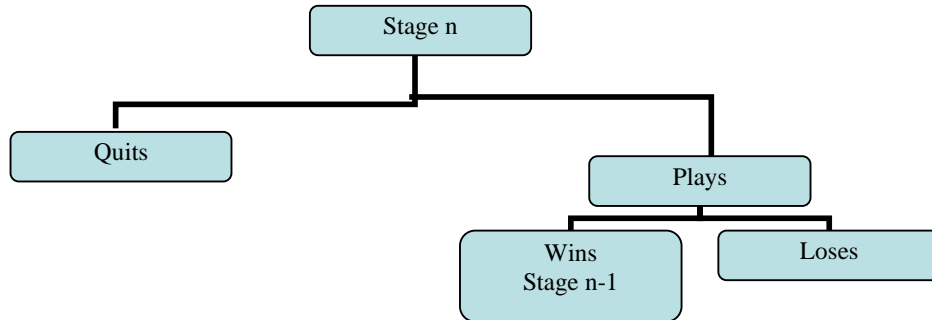
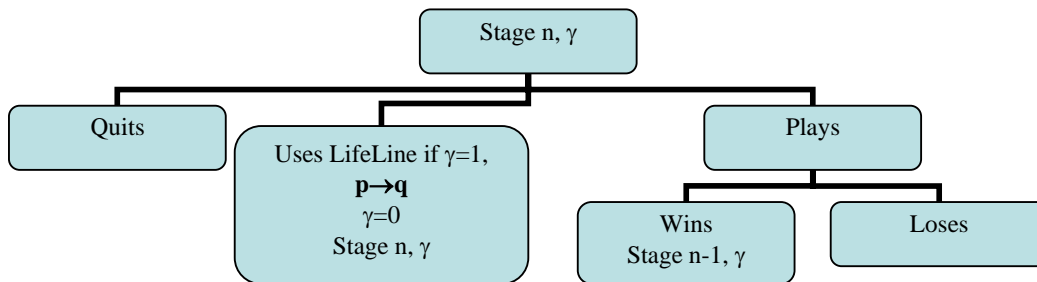


Figure 1b With Life Line



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, if $\mathbf{p} = (\frac{1}{4}, \frac{1}{4}, \frac{1}{4}, \frac{1}{4})$, by symmetry the possible realisations of \mathbf{q} after using “50:50” are $\mathbf{q} = (\frac{1}{2}, \frac{1}{2}, 0, 0)$. A contestant who would reject such a “50:50” gamble would place no value on the lifeline. Except in these circumstances, the life line will be used if $a_{n+1} \leq k_n(\mathbf{p})$, otherwise the contestant will answer and retain the lifeline for future use.

5.4.b. The complete game

We now assume that the three life lines are available but each can be played at most once. As above each life line generates a new belief \mathbf{q} which is used in the

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 life line, 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 $\underline{\gamma} = (\gamma_1, \gamma_2, \gamma_3)$ for the “lifeline state” vector where $\gamma_i = 0$ if the i 'th lifeline has been played and 1 otherwise. Then, $W_n(\mathbf{p}; \underline{\gamma})$ denotes 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:

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

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

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

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

$$k_n(\mathbf{p}; \underline{\gamma}) = \max_{\mathbf{e} \in S(\underline{\gamma})} \{ \mathbb{E}[W_n(\mathbf{Q}_e; \underline{\gamma} - \mathbf{e}) | \mathbf{p}] \}, \quad (28)$$

where $S((\gamma_1, \gamma_2, \gamma_3)) = \{ \mathbf{e} \in R^3 : e_i \in \{0, \gamma_i\} \text{ for } i = 1, 2, 3 \text{ and } \mathbf{e} \neq \mathbf{0} \}$ and \mathbf{Q}_e is the distribution of the belief vector when \mathbf{e} is the indicator vector of the lifeline chosen. Note that $S((\gamma_1, \gamma_2, \gamma_3))$ has $\gamma_1 + \gamma_2 + \gamma_3$ elements all of which are unit vectors. If the r 'th unit vector achieves the maximum in (a), then the contestant does best to choose the r 'th lifeline and the distribution of \mathbf{Q}_e is $H_r(\mathbf{q}; \mathbf{p})$. This formulation does not preclude an individual from using more than one lifeline on the same question, a behaviour we observe in some contestants.

5.4.c. “50:50”

This is the simplest life line to model. It provides the candidate with “perfect information” since two incorrect answers are removed. Ex-ante (i.e. before the life line is played) the contestant believes that the correct answer is i ($=1, \dots, 4$) 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

$$\mathbf{q}_k^{\{i,j\}} = \begin{cases} \frac{p_i}{p_i + p_j} & \text{if } k = i, \\ \frac{p_j}{p_i + p_j} & \text{if } k = j \\ 0 & \text{otherwise.} \end{cases} \quad (29)$$

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

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

5.4.d. “Ask the Audience”

Modeling the “Ask the Audience” life line requires more than simply applying Bayes’ rule to the current belief draw. In particular we must allow the candidate to learn from the information provided by the life line, 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 \mathbf{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, $\Pr[\mathbf{inf}]$, and the individual’s

belief about the likelihood that the information will arise if A subsequently occurs, $\Pr[\text{inf} | A]$ according to the following rule, related to Bayes theorem:

$$\Pr[A | \text{inf}] = \Pr[\text{inf} | A] \Pr[A] / \Pr[\text{inf}]. \quad (30)$$

In this expression $\Pr[\text{inf} | A]$ is understood as another component of the individual's belief, his/her assessment of the likelihood of the signal given that the relevant event subsequently occurs.

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

$$\Pr[A | \text{inf}] = \frac{\Pr[\text{inf} | A] \Pr[A]}{\Pr[\text{inf} | A] \Pr[A] + \Pr[\text{inf} | \bar{A}] \Pr[\bar{A}]}. \quad (31)$$

In our context we understand the appeal to the audience as an appeal to an expert, and assume that the events of interests 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

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

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 \mathbf{q} is:

$$\pi_k = \theta_k p_k / \sum_{j=1}^4 \theta_j p_j. \quad (33)$$

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 \mathbf{p} . Given \mathbf{p} , the density of the updated belief $H_2(\pi; \mathbf{p})$ can be calculated as:

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

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

The quantities $\Pr[\mathbf{q} = (q_1, q_2, q_3, q_4) | \text{answer } k \text{ is correct}] \equiv \theta_k$ represent the added information obtained from using the life line and are estimable from the data provided we assume a form of conditional independence. In particular we require that the candidate's choice to ask the audience does not influence the audience's answer.

Furthermore, our assumptions concerning the generation of the questions imply that there is no information contained in the position of the correct answer, hence we expect the following symmetry restrictions to hold :

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

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 $\Pr[\mathbf{q} = (q_1, q_2, q_3, q_4) | \text{answer } k \text{ is correct}]$.

In practice we assume that, given answer k is correct, information \mathbf{q} has a Dirichlet density $D(\mathbf{q}; \gamma_k(\lambda, \nu))$, $k=1\dots 4$, defined over Δ_4 such that

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

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 life line is used, and the observation of the correct answer. For completeness note that θ_k can be defined in

terms of the elements of \mathbf{q} as $\theta_k = q_k^{\lambda-\nu} / \sum_{j=1}^4 q_j^{\lambda-\nu}$. The information density which the

candidate 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:

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

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

5.4.e. "Phone a Friend"

To use this lifeline the candidate determines ahead of the game six potential experts, "friends", and when she plays the life line she chooses one in this list of six. Obviously one imagines that the candidate engages in some diversification when drawing 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 the choice of a particular friend the candidate is likely to assign the expert optimally. There is however little information available to us about this process. As a consequence our model for this particular life line 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 candidate of his confidence. Hence either the candidate knows the answer and his/her opinion "swamps" the candidate's belief, or the expert is ignorant and conveys no information and the candidate's belief is left unchanged. The density of the updated belief is therefore:

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

6. Econometric specification and estimation

6.1. Specification of the belief distribution

The distribution of the beliefs is one of the main element of the model since it describes the distribution of the unobservables. Under the assumptions we make above (see section 5.2) the joint density $\psi_4(\cdot)$ 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, $B(\alpha, \alpha)$ ⁷, i.e.:

$$\phi(x) = \frac{\Gamma(2\alpha)}{\Gamma(\alpha)^2} x^{\alpha-1} (1-x)^{\alpha-1} \mathbf{1}_{[x \in [0,1]]}, \quad (37)$$

with α some positive parameter, and where $\Gamma(u)$ is the gamma function. For any random variable following a symmetric beta distribution the expectation is $\frac{1}{2}$ and

$$\mu_2 = \frac{\Gamma(2\alpha)\Gamma(\alpha+2)}{\Gamma(\alpha)\Gamma(2\alpha+2)} = \frac{1}{2} \left(\frac{\alpha+1}{2\alpha+1} \right).$$

In what follows, it will prove necessary to obtain ordered draws from the joint distribution of ordered statistics of the belief distribution. Because of the symmetry assumptions that we impose on $\psi_4(\cdot)$, the joint density function of the order statistics (i.e. the vector of beliefs ordered in decreasing order) is simply $4!\psi_4(\tilde{\mathbf{p}})$, where $\tilde{\mathbf{p}}$ is a vector of values ordered in decreasing order. From the definition of $\psi_4(\cdot)$ note that it is a mixture with equal weights of $4!=24$ densities of the form:

$$\chi_4(x_1, x_2, x_3, x_4) = \frac{2}{\mu_2} \phi(x_1) \phi\left(\frac{x_2}{1-x_1}\right) \phi\left(\frac{x_3}{1-x_1-x_2}\right), \quad (38)$$

and the mixture is taken over all permutations of the argument. Note however that $\psi_4(\cdot)$ and $\chi_4(\cdot)$ share the same density for the order statistics. Clearly any $4!$ permutations of any given draw will lead to the same order statistics. Hence a given draw $\mathbf{x} = (x_1, x_2, x_3, x_4)$ from the order statistics for $\chi_4(\cdot)$ occurs with probability:

$$\frac{2}{\mu_2} \sum_{\{i,j,k,l\} \in P_4} \phi(x_i) \phi\left(\frac{x_j}{1-x_i}\right) \phi\left(\frac{x_k}{1-x_i-x_j}\right), \quad (39)$$

and this is exactly equal to $4!\psi_4(\mathbf{x})$.

Since it is straightforward to rank four numbers in decreasing order, the last issue is to draw from a multidimensional random variable with joint density $\chi_4(\cdot)$.

⁷ The density of a random variable following a general Beta distribution is :

$$\beta(z; a, b) = \frac{\Gamma(a+b)}{\Gamma(a)\Gamma(b)} (1-z)^{b-1} (z)^{a-1} \mathbf{1}_{[z \in [0,1]]}, \text{ with } a, b > 0$$

The result shown in Appendix 1 shows how this can be done straightforwardly since $\chi_4(\cdot)$ can be factorised as follows

$$\begin{aligned} & \chi_4(x_1, x_2, x_3, x_4) \\ &= \beta(x_1; a, a+2) \frac{1}{1-x_1} \beta\left(\frac{x_2}{1-x_1}; a, a+1\right) \frac{1}{1-x_1-x_2} \beta\left(\frac{x_3}{1-x_1-x_2}; a, a\right), \end{aligned}$$

which suggest that a draw from $\chi_4(\cdot)$ can be obtained from three independent draws from distinct Beta distributions with parameters $(a, a+2)$, $(a, a+1)$ and (a, a) respectively.

6.2. 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 online⁸

6.2.a. Calculating the probabilities when only one life line is available.

When the candidate has used all his/her life lines before stage n , the events of interest are the occurrences of the candidate quitting or losing, and for the last question the event that the candidate 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 than one) life line is available the calculations are made more complicated because of the information which is gained when the lifeline is used and which allows the candidate to update their belief. Hence given the initial draw of the belief we are required to define whether this particular draw leads to the use of the (a) life line and then whether the updated belief (if the life line is played), or the original belief (if the life line is not played) is informative enough to lead the candidate 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 is observed (which is defined precisely below) given that the candidate is in the life line 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. Moreover, $\Omega_{k,n}^{ijk}$ is the expected value

⁸ <http://www.qub.ac.uk/schools/SchoolofManagementandEconomics/Staff/LanotGauthier/>

of $\Omega_{k,n}^{ijk}(\mathbf{p})$ over all possible values of realisations of \mathbf{p} , i.e. $\Omega_{k,n}^{ijk} = \mathbb{E}[\Omega_{k,n}^{ijk}(\mathbf{P})]$.

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 (which is defined precisely below) given that the candidate starts the question in the life line state ijk and transit to life line state $i'j'k'$. We consider below representative events for each life line.

6.2.a.i) "50:50" is the only life line available at stage n .

The candidate uses "50:50", plays and wins (moves to the next stage or wins the million prize).

Define first the probability for the candidate to use "50:50", play and win given a draw (ordered in decreasing order) \mathbf{p} from the belief :

$$\Pr[\{\text{use "50:50"}\} \wedge \{\text{plays}\} \wedge \{\text{wins}\} | \{\text{stage } n\}, \mathbf{p}] \equiv \Omega_{1,n}^{100}(\mathbf{p}) = \mathbf{1}_{\left[k_n^1(\mathbf{p}, 0, 0, 0) \geq p_1(f_{n-1}(1, 0, 0) - b_n) + b_n \right]} \Omega_{1,n}^{100,000}(\mathbf{p}), \quad (40)$$

where $\Omega_{1,n}^{100,000}(\mathbf{p}) = \frac{1}{3} \sum_{j=1}^3 p_j \sum_{k=j+1}^4 \mathbf{1}_{\left[\pi_{jk}(\mathbf{p})(f_{n-1}(0, 0, 0) - b_n) + b_n \geq a_{n+1} \right]}$, with $\pi_{jk} = \frac{p_j}{p_j + p_k}$. This last

expression is the probability that given \mathbf{p} the candidate answers correctly after using the life line, i.e. wins. Hence the unconditional probability is such that

$$\Pr[\{\text{use "50:50"}\} \wedge \{\text{plays}\} \wedge \{\text{wins}\} | \{\text{stage } n\}] \equiv \Omega_{1,n}^{100} = \int_{\tilde{\Delta}_4} \Omega_{1,n}^{100}(\mathbf{p}) \tilde{\psi}_4(\mathbf{p}) d\mathbf{p} \quad (41)$$

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

where $\mathbf{p} = (p_1, p_2, p_3, p_4)$, $\tilde{\Delta}_4$ 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 candidate with a life line available will either use it (or perhaps then quit) or play. It is then straightforward to verify that these five expressions above sum to unity, in particular the sum of the first three expression sum to the probability that the candidate uses the life line, i.e. the complement to the sum of the last two probabilities.

Each term of the sum that determine $\Omega_1^{100}(\mathbf{p})$ (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 life line is played, with probability $\frac{1}{3}(p_j + p_k)$, multiplied by the probability that the remaining alternative with the largest updated belief is correct, with probability $\pi_{jk}(\mathbf{p}) = \frac{p_j}{p_j + p_k}$ with $p_j \geq p_k$, multiplied by the indicator that, given the updated belief, the candidate decides to play.

6.2.a.ii) “Ask the Audience” is the only life line left at stage n .

The candidate uses “Ask the Audience”, plays and loses.

$$\Pr[\{\text{use "Ask the Audience"}\} \wedge \{\text{plays}\} \wedge \{\text{loses}\} | \{\text{stage } n\}] \equiv \Omega_{2,n}^{010} = \int_{\tilde{\Delta}_4} \Omega_{2,n}^{010}(\mathbf{p}) \tilde{\psi}_4(\mathbf{p}) d\mathbf{p}, \quad (42)$$

where $\Omega_{2,n}^{010}(\mathbf{p}) = \mathbf{1}_{[k_n^2(\mathbf{p},0,0,0) \geq p_1(f_{n-1}(0,1,0)-b_n)+b_n]} \Omega_{2,n}^{010,000}(\mathbf{p})$, and $\Omega_{2,n}^{010,000}(\mathbf{p}) = \int_{\tilde{\Delta}_4} (1 - \pi_1(\mathbf{q};\mathbf{p})) \mathbf{1}_{[\pi_1(\mathbf{q};\mathbf{p})(f_{n-1}(0,0,0)-b_n)+b_n \geq a_{n+1}]} D(\mathbf{q};\mathbf{p}, \lambda, \nu) d\mathbf{q}$ 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 $\pi(\mathbf{q};\mathbf{p})$.

6.2.a.iii) “Phone a Friend” is the only life line left at stage n .

The candidate uses “Phone a Friend” and quits.

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

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

6.2.b. General Case: all the life lines are available

When more than one life line is available at stage n the number of elementary events of interest increases since not only the candidates can decide to play one life line among many but the candidate can decide to play several life lines to answer any given question. Hence while there is only five elementary events of interest when only one given life line is available there are nine such events when two particular life lines are available and seventeen when the three life lines are available, ignoring the order in which the candidate uses the life-line and not counting events with zero probability ex-ante (for example observing an event such as quitting while the three life lines are

available)⁹. In this 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 candidate uses the three life lines (in any order), plays and loses.

$$\begin{aligned}
& \Pr\left[\{\text{uses all life lines}\} \wedge \{\text{plays}\} \wedge \{\text{loses}\} \mid \{\text{stage } n\}\right] \equiv \Omega_{2,n}^{111} = \\
& \int_{\tilde{\Delta}_v} \\
& \mathbf{1}_{\left[k_n^1(\mathbf{p},0,1,1) \geq \max\{p_1(f_{n-1}(1,1,1)-b_n)+b_n, k_n^2(\mathbf{p},1,0,1), k_n^3(\mathbf{p},1,1,0)\}\right]} \Omega_{2,n}^{111,011}(\mathbf{p}) \\
& \mathbf{1}_{\left[k_n^2(\mathbf{p},1,0,1) \geq \max\{p_1(f_{n-1}(1,1,1)-b_n)+b_n, k_n^1(\mathbf{p},0,1,1), k_n^3(\mathbf{p},1,1,0)\}\right]} \Omega_{2,n}^{111,101}(\mathbf{p}) + \\
& \mathbf{1}_{\left[k_n^3(\mathbf{p},1,1,0) \geq \max\{p_1(f_{n-1}(1,1,1)-b_n)+b_n, k_n^1(\mathbf{p},0,1,1), k_n^2(\mathbf{p},1,0,1)\}\right]} \Omega_{2,n}^{111,110}(\mathbf{p}) \\
& \tilde{\psi}_4(\mathbf{p}) d\mathbf{p}.
\end{aligned} \tag{43}$$

$$\text{where } \Omega_{2,n}^{111,011}(\mathbf{p}) = \frac{1}{3} \sum_{j=1}^3 \sum_{k=j+1}^4 \Omega_{2,n}^{011}(\pi_{j,k}(\mathbf{p}), \tau_{k,j}(\mathbf{p}), 0, 0), \tag{44}$$

$$\Omega_{2,n}^{111,101}(\mathbf{p}) = \int_{\Delta_v} \Omega_{2,n}^{101}(\pi(\mathbf{q}; \mathbf{p})) D(\mathbf{q}; \mathbf{p}, \lambda, \nu) d\mathbf{q}, \tag{45}$$

$$\Omega_{2,n}^{111,110}(\mathbf{p}) = \kappa \Omega_{2,n}^{110}(1, 0, 0, 0) + (1 - \kappa) \Omega_{2,n}^{110}(\mathbf{p}). \tag{46}$$

Inspection of these expressions reveals that the probabilities of events such that more than one life line is available, here $\Omega_{2,n}^{111}$, can be defined recursively in terms of the conditional probability of events with one life line less, 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 life line, i.e. $\Omega_{2,n}^{001}(\mathbf{p})$, $\Omega_{2,n}^{100}(\mathbf{p})$ and $\Omega_{2,n}^{010}(\mathbf{p})$. This property is clearly a consequence of the recursive definition of the value function over the life-line part of the state space (see section 5.4.b).

⁹ In the case of two life lines left : 1) Uses the two life lines, plays and wins; 2) Uses the two life lines, plays and loses; 3) Uses the two life lines, plays and loses; 4) Uses one of two life lines, plays and wins; 5) Uses one of two life lines, plays and loses; 6) Uses other life line, plays and wins; 7) Uses other life line, plays and loses; 8) Does not use any life line, play and win; 9) Does not use any life line, play and loses;

In the case of three life line left: 1) Uses the three life lines, plays and wins; 2) Uses the three life lines, plays and loses; 3) Uses the three life lines, 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 life line, play and win; 17) Does not use any life line, play and loses;

Recall, however, that the number of events of interest when the three life lines are available is larger than when only two or less are available. Hence the definition of 17 probabilities with three life line at stage n , i.e. $\Omega_{m,n}^{111}$, $m=1\dots 17$, will involve the 27 conditional probabilities with two life lines, i.e. $\Omega_{m,n}^{011}(\mathbf{p})$, $\Omega_{m,n}^{101}(\mathbf{p})$ and $\Omega_{m,n}^{110}(\mathbf{p})$, $m=1\dots 9$. In turn each of these conditional probabilities will depend on the 15 probabilities with one life line as defined in the previous section, i.e. $\Omega_{m,n}^{100}(\mathbf{p})$, $\Omega_{m,n}^{010}(\mathbf{p})$ and $\Omega_{m,n}^{001}(\mathbf{p})$ $m=1\dots 5$.

The three life lines are available, the candidate uses “50:50”, plays and wins.

$$\begin{aligned} & \Pr[\{\text{uses "50:50" only among 3 life lines}\} \wedge \{\text{plays}\} \wedge \{\text{wins}\} | \{\text{stage } n\}] \equiv \Omega_{10,n}^{111} \\ & = \int_{\tilde{\Delta}_4} \mathbf{1}_{\left[k_n^1(\mathbf{p},0,1,1) \geq \max\{p_1(f_{n-1}(1,1,1)-b_n)+b_n, k_n^2(\mathbf{p},1,0,1), k_n^3(\mathbf{p},1,1,0)\} \right]} \Omega_{10,n}^{111,011}(\mathbf{p}) \tilde{\psi}_4(\mathbf{p}) d\mathbf{p}. \end{aligned}$$

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

Three life lines are available, the candidate does not use any, plays and loses.

$$\begin{aligned} & \Pr[\{\text{does not use any of the 3 life lines}\} \wedge \{\text{plays}\} \wedge \{\text{loses}\} | \{\text{stage } n\}] \equiv \Omega_{17,n}^{111} \\ & = \int_{\tilde{\Delta}_4} \mathbf{1}_{\left[p_1(f_{n-1}(1,1,1)-b_n)+b_n \geq \max\{k_n^1(\mathbf{p},0,1,1), k_n^2(\mathbf{p},1,0,1), k_n^3(\mathbf{p},1,1,0)\} \right]} (1-p_1) \tilde{\psi}_4(\mathbf{p}) d\mathbf{p}. \end{aligned}$$

6.3. Simulation and smoothing

The evaluation of the probabilities $\Omega_{m,n}^{rst}(\mathbf{p})$, $n=1..15$, $m=1..17$ ¹⁰, $(r,s,t) \in \{0,1\}^3$ 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. Under the specification of the belief we describe above simulation methods (as described in Gouriéroux and Monfort (1996) and Train (2003)) that 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 perfectly to a simulation based likelihood methodology since simulations of Beta variates are obtained simply

¹⁰ If $\Omega_{m,n}^{rst}$ is not defined for some m , and some r,s,t we assume $\Omega_{m,n}^{rst} = 0$.

from Gamma variates (see for example Poirier (xxxx)). In turn Gamma variates themselves can be obtained 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)¹¹. The main advantage of their results is that it allows for simulations that are continuous in the parameters of the Gamma distributions. Evaluation by simulations 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

$$\Omega_{1,n}^{100} = \int_{\Delta_4} \Omega_{1,n}^{100}(\mathbf{p}) \tilde{\psi}_4(\mathbf{p}) d\mathbf{p} = \int_{\Delta_4} \mathbf{1}_{\left[k_n^1(\mathbf{p}, 0, 0, 0) \geq p_1(f_{n-1}(1, 0, 0) - b_n) + b_n \right]} \Omega_{1,n}^{100,000}(\mathbf{p}) \tilde{\psi}_4(\mathbf{p}) d\mathbf{p}, \quad (47)$$

can be approximated by

$$\hat{\Omega}_{1,n}^{100}(S) = \frac{1}{S} \sum_{s=1}^S \Omega_{1,n}^{100}(\mathbf{p}_s) = \frac{1}{S} \sum_{s=1}^S \mathbf{1}_{\left[k_n^1(\mathbf{p}_s, 0, 0, 0) \geq p_{1,s}(f_{n-1}(1, 0, 0) - b_n) + b_n \right]} \Omega_{1,n}^{100,000}(\mathbf{p}_s), \quad (48)$$

where \mathbf{p}_s is one of S (the number of simulations) independent draws from the distribution of the order statistics of the belief, $\tilde{\psi}_4(\cdot)$. In fact the accuracy of this simulated probability (and of all others which involve draws from $\tilde{\psi}_4(\cdot)$) can be improved upon through antithetic variance reduction techniques which involve the permutations of the gamma variates used to generate each individual beta variate¹² (as explained for example in Davidson and McKinnon (2004) or in Train (2003)). Moreover the quantity

$$\Omega_{10,n}^{111} = \int_{\Delta_4} \mathbf{1}_{\left[k_n^1(\mathbf{p}, 0, 1, 1) \geq \max\{p_1(f_{n-1}(1, 1, 1) - b_n) + b_n, k_n^2(\mathbf{p}, 1, 0, 1), k_n^3(\mathbf{p}, 1, 1, 0)\} \right]} \Omega_{10,n}^{111,011}(\mathbf{p}) \tilde{\psi}_4(\mathbf{p}) d\mathbf{p}. \quad (49)$$

can be evaluated simply by

$$\hat{\Omega}_{10,n}^{111}(S) = \frac{1}{S} \sum_{s=1}^S \mathbf{1}_{\left[k_n^1(\mathbf{p}_s, 0, 1, 1) \geq \max\{p_{1,s}(f_{n-1}(1, 1, 1) - b_n) + b_n, k_n^2(\mathbf{p}_s, 1, 0, 1), k_n^3(\mathbf{p}_s, 1, 1, 0)\} \right]} \Omega_{10,n}^{111,011}(\mathbf{p}_s), \quad (50)$$

¹¹ This is implemented in Gauss in the procedure **gammai** (contained in the file cdfchic.src)

¹² 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 $\alpha > 0$).

or any improvement of it. Similarly $\Omega_{2,n}^{111,101}(\mathbf{p}) = \int_{\Delta_4} \Omega_{2,n}^{101}(\pi(\mathbf{q};\mathbf{p}))D(\mathbf{q};\mathbf{p},\lambda,\nu)d\mathbf{q}$ can be evaluated by $\hat{\Omega}_{2,n}^{111,101}(\mathbf{p};S) = \frac{1}{S} \sum_{i=1}^4 p_i \sum_{s=1}^S [\Omega_{2,n}^{101}(\pi(\mathbf{q}_{s,i};\mathbf{p}))]$, where $\mathbf{q}_{s,i}$ is one of S independent draws from $D(\mathbf{q};\gamma_i(\lambda,\nu))$.

Finally all quantities $k_n^2(\mathbf{p},r,s,t) \equiv E_{\pi_2|\mathbf{p}}[W_n(\Pi,r,s,t)|\mathbf{p}]$ 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 $\hat{k}_n^2(\mathbf{p},r,s,t;S) = \frac{1}{S} \sum_{i=1}^4 p_i \sum_{s=1}^S W_n(\mathbf{q}_{s,i},r,s,t)$, where $\mathbf{q}_{s,i}$ is one of S independent draws from $D(\mathbf{q};\gamma_i(\lambda,\nu))$.

In practice these expression are modified in order to smooth out the discontinuities that are created by the indicators terms. Hence the terms $\mathbf{1}_{[v_1 \geq \max\{v_2, v_3, v_4\}]}$, $\mathbf{1}_{[v_1 \geq \max\{v_2, v_3\}]}$, or $\mathbf{1}_{[v_1 \geq v_2]}$, are replaced by smoothed versions, respectively,

$$\frac{1}{1 + \exp(\eta(v_2 - v_1)) + \exp(\eta(v_3 - v_1)) + \exp(\eta(v_4 - v_1))}, \frac{1}{1 + \exp(\eta(v_2 - v_1)) + \exp(\eta(v_3 - v_1))}$$

, and $\frac{1}{1 + \exp(\eta(v_2 - v_1))}$ where η is a smoothing constant. In the limit as $\eta \rightarrow +\infty$

the smoothed versions tend to the indicators.

6.4. Likelihood.

The contribution to the likelihood for some individual history is the product of the probabilities of success and of the particular pattern of use for the life lines for that individual history up to and including the penultimate question, multiplied by the probability that for his/her last question the candidate wins a million, loses or quits and the observed use of the life-lines for this last question.

We assume that the expected utility function takes the form $U(c) = (c + \gamma)^{1-\rho} / (1-\rho)$ which features the CRRA assumption and treats initial wealth, γ , as a parameter to be estimated¹³.

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

¹³ Later we also consider a generalisation of the CRRA function that encompasses both CRRA and CARA.

$$\hat{L}_{S,i} \left(\left\{ (LL(k,i), ll(k,i)) \right\}_{k=1}^{n_i^*}; (\alpha, \rho, \gamma, \kappa); (\hat{\lambda}, \hat{\nu}) \right) = \left\{ \prod_{k=1}^{n_i^*-1} \hat{\Omega}_{ll(k,i),k}^{LL(k,i)} \right\} \hat{\Omega}_{ll(n_i^*,i),n_i^*}^{LL(n_i^*,i)}, \quad (51)$$

where $LL(k,i)$ indicates the number and nature of the life lines available to the candidate i at stage k , and $ll(k,i)$ selects the relevant probability depending on the life line used by candidate i at stage k . $(\alpha, \rho, \gamma, \kappa)$ is the vector of parameters of interest, i.e. α is the parameters 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 unknown parameter in the distribution of the updated belief which results from the use of “phone a friend”. Finally $(\hat{\lambda}, \hat{\nu})$ are the independent estimates of the parameters of the density of the updated belief which results from the use of “Ask the Audience”.

7. The Data

For each broadcast show the operator, Celador PLC, selected 10 names at random from a (large) list of entrants 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 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 to June 2003. This gives us information on the behaviour of 515 contestants¹⁴ who played the second round sequence of multiple choice questions.

¹⁴ 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.

However, a major concern about the findings of the gameshow literature is that the data is generated by selected samples¹⁵. 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 players 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 population¹⁶.

Questionnaire replies were received by 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 players. Not surprisingly, these second round players 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 – something that is consistent with previous papers which have found that men to be less risk averse than women. 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 one very clear difference between WWTBAM entrants and the population sample survey data is that they are much more likely to be male. Attempting to enter the gameshow is risky and this would be consistent with the finding, in some of the earlier literature, that women are more risk averse. Of course, it would be consistent with other hypotheses too and no specific inferences can be made from this gender difference. The table also shows the corresponding information from various social surveys, re-weighted to match the gender mix in the questionnaire data.

¹⁵ In fact, Hersch and McDougall (1977) and Fullenkamp *et al* (2003) do report some comparisons between players and the population and finds 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 players do seem to have different characteristics than non-players (see Farrell and Walker (1999)).

¹⁶ 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

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.

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 with the population of 515 actual contestants. We have no consistent information on the characteristics of players in the

population apart from what we see on screen. Thus, Table 2 records on 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 differences are not significant. Thus, we can have some confidence that the representativeness of players (in the questionnaire data) carries over to the population data in the videotapes.

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.

The distribution of winnings, for the second round contestants, depends on whether the player quit or failed to answer the last question asked. Figure 1 illustrates using the videotape data where a small amount of jitter has been added to the data to show the nature of the joint distribution of quitting and stage of the game. Almost all players who survived beyond £125,000 quit rather than failed – only one player failed at £500,000 and so went away with just £32,000 instead of quitting and going away with £250,000. Those that failed to answer correctly the last question that they were asked went away with their corresponding value of b , the reserve level of winnings. Only three contestants failed at a sub £1000 question and went away with nothing. Three players won the £1m prize. Two-thirds of players 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)¹⁷.

Finally, the use of lifelines is an important part of observed behaviour that our model attempts to explain. 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.

¹⁷ Here we categorise those that won the maximum £1m as quitters.

Figure 1 Distribution of winnings (% £ ,000)

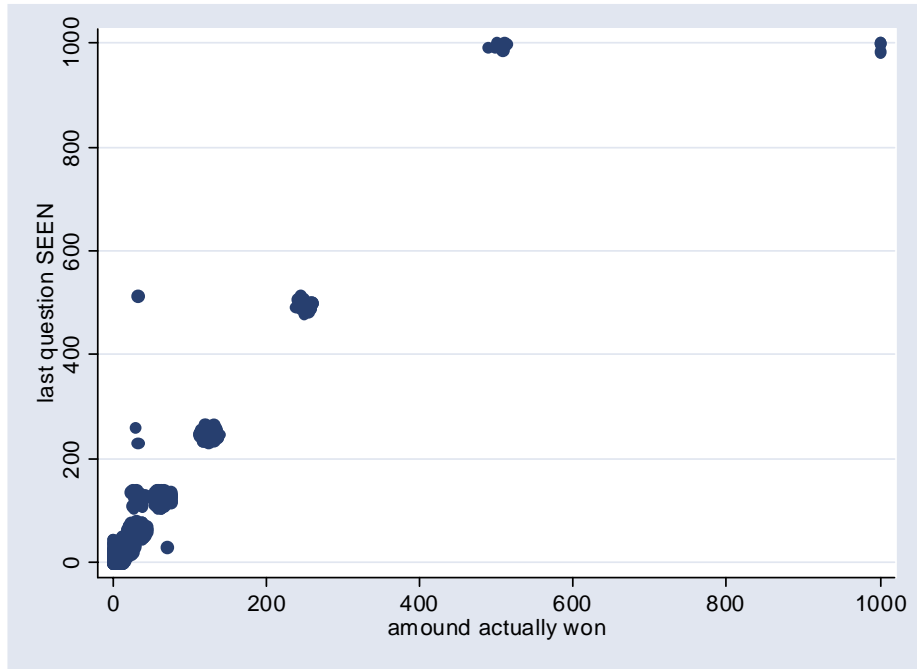
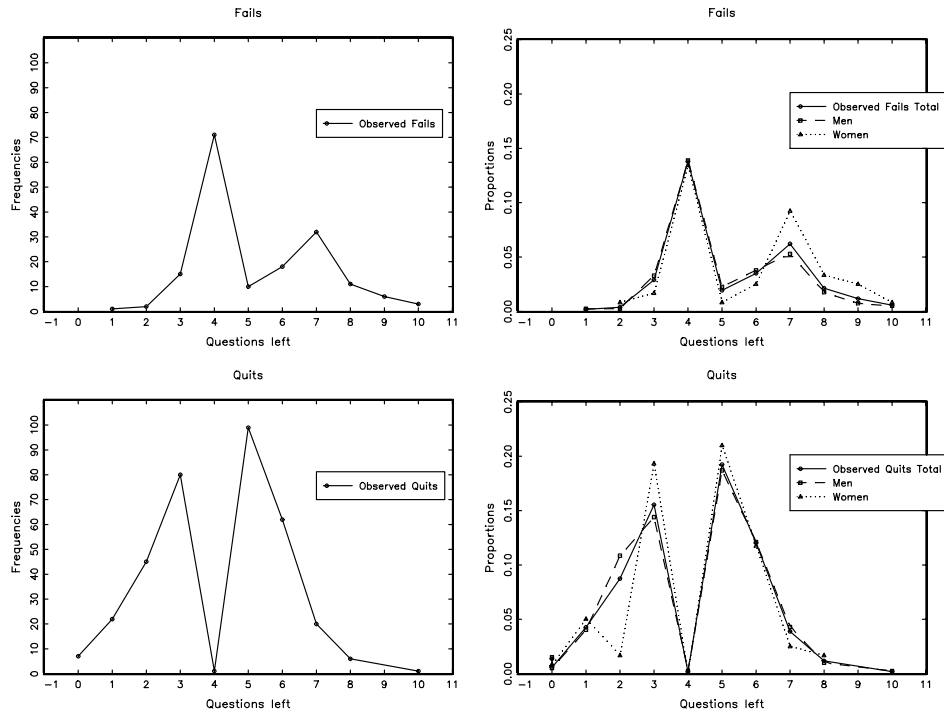


Figure 2 Observed Fails and Quits Frequencies and Rates



8. Estimation and Results

To estimate the parameters of the model, we first estimate our model for the histograms that are produced by the life line “Ask the Audience” using the data we have collected on this histogram and on our knowledge of what the correct answer was. The parameters estimates are presented in Table 3. Assume that the first candidate answer is the correct one, these estimates imply that on average we expect the life line “Ask the Audience” to produce the histogram (0.63, 0.12, 0.12, 0.12). These parameters allow us to evaluate the quality of the lifeline in the manner described in section 5.4.d above.

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	

Treating these parameters as constants we then proceed to estimate the remaining parameters of the model. Table 4 presents the preference parameters as well as the estimate of the probability that the chosen friend, in the PAF lifeline, knows the correct answer is $\kappa \approx 0.41$.

Our preferred estimated value for the coefficient of relative risk aversion is remarkably close to 1 (although statistically significantly different from 1). The parameter γ , which can be interpreted as initial wealth measured in thousands of pounds, is significantly estimated at 0.41 (i.e. an initial “wealth” of £410).

Two additional parameters which allow for the distribution of the initial belief to change with each round of the game are estimated. To illustrate how the distribution of the beliefs changes as the game progresses we have calculated, in Figure 3, the distribution of the maximum belief when one (respectively 3, 5, 8, 10) question(s) remain to be played. We can contrast this distribution at the beginning of the game, where the maximum belief is relatively concentrated around 1, to the later

rounds of the game, where the maximum belief is in fact concentrated away from 1. To see this compare the slopes of the distribution functions to the left of 1 - in the former case the slope is large while in the latter case the slope is close to zero.

Table 4 Maximum Likelihood Estimates

Parameters	Estimates	Std. err.
ρ	1.018	0.001
γ	0.410	0.077
$-\ln\left(\frac{1}{\kappa}-1\right)$	-0.325	0.112
Number of observations	515	-
Log-Likelihood	-4543.8141	-

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

Figure 4 describes the value of playing the game as a function of the number of questions remaining (on the x-axis) and the number and the 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. “Ask the Audience” appears to be the most valuable lifeline while “5050” and “Phone a Friend” have almost identical values. In fact, the model predicts that “Ask the Audience” is almost as valuable as “5050” and “Phone a Friend” together.

In Figure 5 we use the estimates to compute the predicted probabilities of quitting and failing at each question 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.

Finally, in Table 5, we compute the certainty equivalent of the gambles taken at each stage of the game for three different types of individual. The top third of the table corresponds to a very able player, the middle third is about a typical individual, while the bottom third is for a low ability player. The certainty equivalents measure in money terms the value of being able to play the game and take into account the value

of being able to play further if the player is successful at the current stage. Moreover we present similar calculations for the value of playing the lifeline (again given a particular draw). Clearly the belief has a substantial effect on the certainty equivalents. Indeed our model predicts that if faced with either of the second or third belief draw, candidates would be prepared to pay sizeable amounts (amounts larger than £300,000 in the case of the million pounds question) to avoid having to answer the question. Lifelines are clearly valuable when the belief draw is not an extreme one.

Figure 3

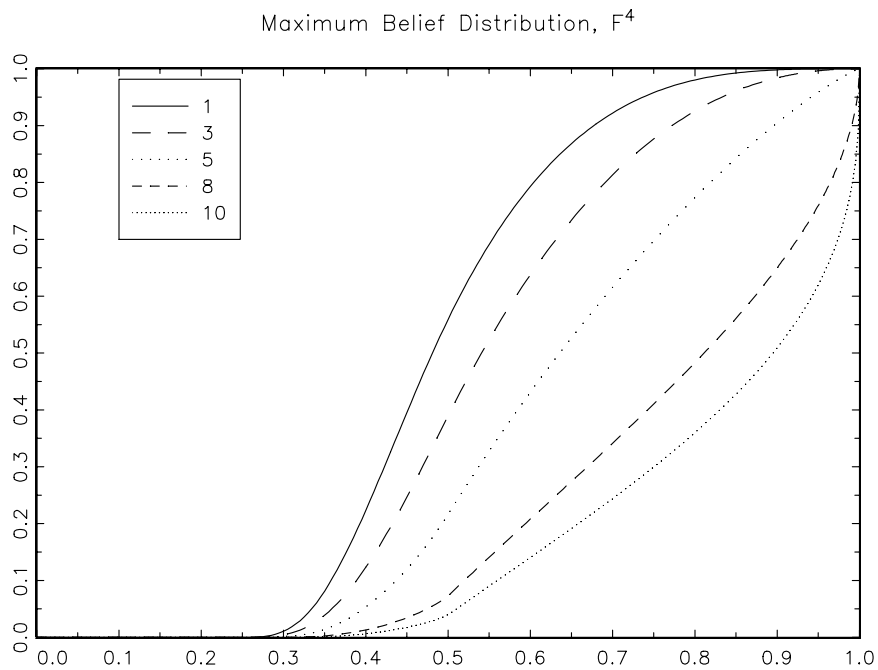


Figure 4 Value of playing the game at stage n , given the life line state.

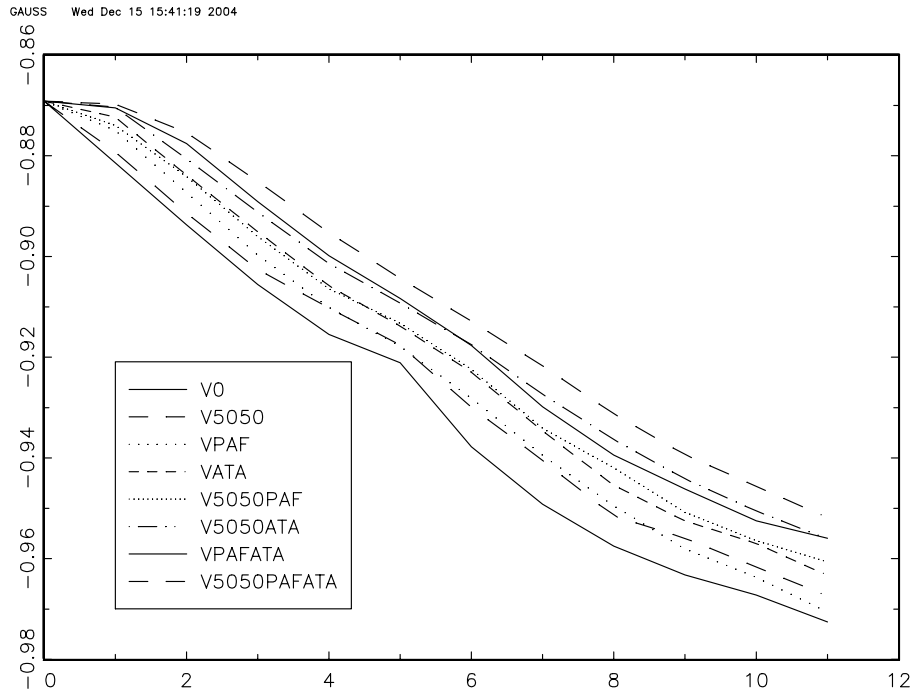


Figure 5 Observed versus Predicted Frequencies of Fails and Quits

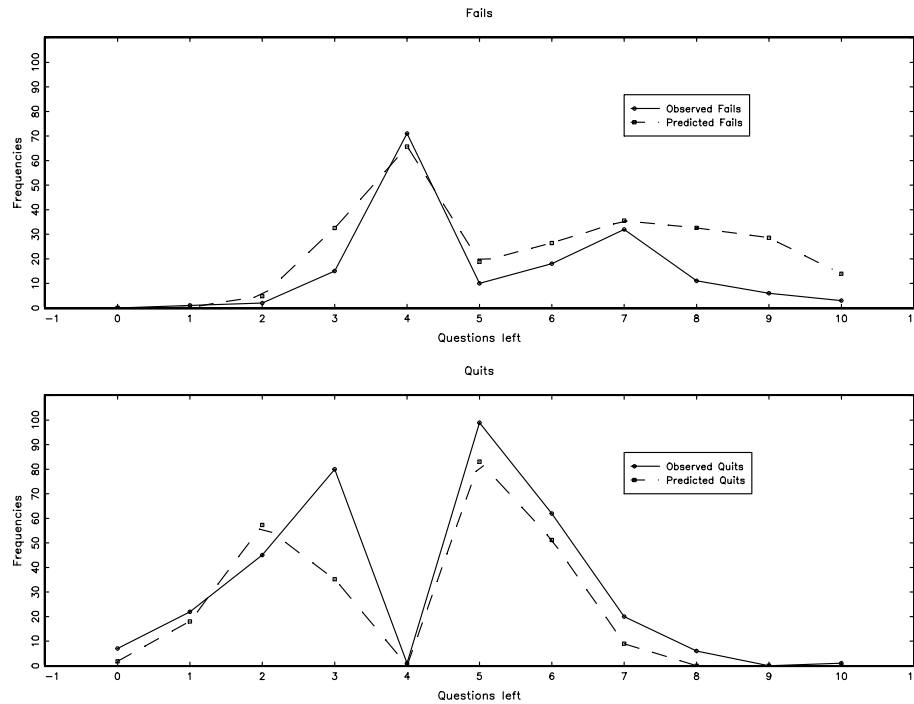


Table 5 Certainty Equivalents (£,000)

		5050		PAF		ATA	
		CE ^a	CE of LL ^b	CE ^a	CE of LL ^b	CE ^a	CE of LL ^b
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

^a: $u^{-1}(p(f_{n-1} - b_n) + b_n)$

^b: $u^{-1}(k_n(\mathbf{p}, \dots))$

9. Conclusions and Extensions

This paper provides new evidence about the degree of individual risk aversion. The analysis is firmly embedded in the expected utility paradigm. Surprisingly, we find the model is broadly effective in explaining behaviour in this simple, and popular, gameshow - *Who Wants to be a Millionaire?* A feature of our analysis is that it is based on data that appears to be representative of the UK population, both in terms of observable characteristics and in terms of other aspects of risk-taking behaviour. Our headline result is that expected utility is approximately logarithmic - the CRRA is 1, with a high degree of precision.

We also use our data to estimate the value of additional information to players in this game of skill. Our headline result is consistent with the results of recent work on the *Hooster Millionaire* gameshow which is the only other game which features, like WWTBAM, such large stakes and involves no complex probability calculations by players.

In part, the paper addresses the challenge to expected utility made by Rabin (2000) who suggests that, since individuals are risk averse when faced with small gambles the implied behaviour with respect to large gambles would be perverse. We indeed find that, in this model with constant CRRA across the huge range of stakes in the game, we do underpredict the extent to which individuals take risk when the stakes are low.

A deficiency of the current work is that we assume that risk aversion does not vary across individuals. We view this as an approximation and we plan to conduct further work that relaxes this by exploiting our questionnaire data which is rich in information about the individuals who played this game, and has been used in this paper only to confirm the representativeness of the videotape data. In particular, we wish to explore the extent to which risk aversion varies with observable characteristics and whether unobserved heterogeneity in risk aversion is correlated with observable risk-taking behaviour.

References

- Adda, J. and R. Cooper (2003), **Dynamic Economics**, MIT Press: Cambridge, Mass.
- Atkinson, A.B. (1977), "Optimal Taxation and the Direct versus Indirect Tax Controversy", **Canadian Journal of Economics**, 10, 590-606.
- 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 (Supplement to 395), 59-73.
- Barsky, R.B. *et al* (1997), "Preference Parameters and Behavioural Heterogeneity: An Experimental Approach in the Health and Retirement Study", **Quarterly Journal of Economics**, 112(2): 537-79
- 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
- Camerer, C. F. (1989), "An Experimental Test of Several Generalized Utility Theories", **Journal of Risk and Uncertainty**, 2(1), 61-104.
- Chetty, R. (2003), "A New Method of Estimating Risk Aversion", **NBER Working Paper** 9988.
- Cicchetti, 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
- 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(2), 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. (1995), "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.
- 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

Proposition (factorisation of $\chi_4(x_1, x_2, x_3, x_4)$):

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

with (x_1, x_2, x_3, x_4) such that $\sum_{i=1}^4 x_i = 1$, $x_i \geq 0$ 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 $\chi_4(\cdot)$.

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

$P = (U_1 \quad \bar{U}_1 U_2 \quad \bar{U}_1 \bar{U}_2 U_3 \quad \bar{U}_1 \bar{U}_2 \bar{U}_3)$, with $\bar{U}_i = 1 - U_i$ for all $i=1..3$, is distributed

with joint density: $\chi_4(x_1, x_2, x_3, x_4) = \frac{2}{\mu_2} \phi(x_1) \phi\left(\frac{x_2}{1-x_1}\right) \phi\left(\frac{x_3}{1-x_1-x_2}\right)$. 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 $\chi_4(\cdot)$ and then to sort the resulting

vector in descending order.