THE UNIVERSITY *of York*

Discussion Papers in Economics

No. 2007/03

Naïve, Resolute or Sophisticated? A Study of Dynamic Decision Making

By

John D Hey and Gianna Lotito

Department of Economics and Related Studies
University of York
Heslington
York, YO10 5DD

Naïve, Resolute or Sophisticated?

A Study of Dynamic Decision Making

John D Hey

LUISS, Rome, Italy and University of York, UK

and

Gianna Lotito

Università del Piemonte Orientale, Alessandria, Italy

Abstract

Dynamically inconsistent decision makers have to decide, implicitly or explicitly, what to do about their dynamic inconsistency. Economic theorists have identified three possible responses – to act naively (thus ignoring the dynamic inconsistency), to act resolutely (not letting their inconsistency affect their behaviour) or to act sophisticatedly (hence taking into account their inconsistency). We use data from a unique experiment (which observes both decisions and evaluations) in order to distinguish these three possibilities. We find that the majority of subjects are either naïve or resolute (with slightly more being naïve) but very few are sophisticated. These results have important implications for predicting the behaviour of people in dynamic situations.

JEL classifications: D90, D80, C91

Keywords: dynamic decision making, naivety, sophistication, resoluteness, dynamic inconsistencies.

Corresponding author:
John D Hey
Department of Economics and Related Studies,
University of York,
Heslington,
York, YO10 5DD,
United Kingdom
jdh1@york.ac.uk

Naïve, Resolute or Sophisticated?

A Study of Dynamic Decision Making

1. Introduction

An important and recurring issue in the analysis of dynamic decision making[1] is the behaviour of dynamically inconsistent people. Do they know that they are dynamically inconsistent, and, if so, what do they do about it? Economic theory has identified three possible responses (though there are obviously many more): that such decision makers act **naively** (ignoring their inconsistency); that they act **resolutely** (not letting their inconsistency affect their behaviour); that they act **sophisticatedly** (taking their inconsistency into account). We report on an experiment that lets us infer which of these responses describes behaviour better. We have designed the experiment in such a way that we can not only observe choices in dynamic decision problems but also we obtain subjects' evaluations of such problems. Combining these two types of data we can estimate the preferences of the decision makers and infer whether they are naïve, resolute or sophisticated.

The problem of inconsistent choice in a dynamic decision problem was initially analysed in the literature in a context of certainty and was related to the problem of preferences changing exogenously through time (see, for example, Strotz 1956). Hammond (1976 and 1977) generalised the analysis, overcoming the distinction between exogenously and endogenously changing tastes, and concentrating the analysis on the essential aspect of the problem – namely that preferences change over time. However he kept it confined to a situation without risk or uncertainty, which will be introduced only later (Hammond 1988a,b; 1989, Raiffa 1968, Machina 1989, McClennen 1990).

The simple example of the drug addict (Hammond 1976) illustrates the problem. Suppose that an individual is considering whether to start taking an addictive drug. The individual[2] would prefer at most to take the drug without consequences. However, he is certain that, if he starts, he will become an addict, with serious consequences for his health. Of course, he can refuse to take the

[1] See, amongst others, Cubitt *et al* (2004), Machina (1989), McClennen (1990).
[2] We presume that he is male to avoid expositional clumsiness.

drug in the first place. This agent is facing a simple dynamic decision problem with the following structure (squares representing decision nodes):

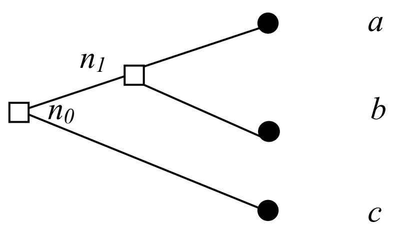

Here three options are available to the agent, which lead to the following outcomes: take the drug till it is harmless, then stop, which leads to outcome *a*; become an addict, which leads to outcome *b*; not take the drug, which leads to outcome *c*.

At the initial decision node n_0 the agent has to decide whether to take the drug or not, and his preferences are $a \succ c \succ b$. If he gets to the choice node n_1 he has become an addict, and therefore the only relevant preferences are those concerning *a* and *b*, and addiction itself means that $b \succ a$. Thus, at n_1 his initial preferences between *a* and *b* get reversed. The agent will choose *b* inconsistently with his previous preference for reaching *a*.

The requirement of *dynamic consistency* is a requirement of consistency between planned choice and actual choice. In the example, the agent decides initially to take the drug, then stop (option *a*), but chooses later not to stop (thus actually choosing option *b*).

Hammond considers *myopic (naïve)* and *sophisticated choice* as two possibilities available to an agent in such a situation of dynamic inconsistency. When acting according to a *myopic approach*, the agent selects at each point those strategies which he judges acceptable from the perspective of that point. In the example, the myopic agent ignores that his tastes are changing, and chooses at each stage the option he considers as best at that moment. Therefore, he will choose option *a* at n_0, but change his mind and choose *b* at n_1. The final outcome will be *b*.

When acting according to a *sophisticated approach*, the agent anticipates his future choice and chooses the best plan among those he is ready to follow to the end: he rejects those plans which imply a choice he anticipates he will not make. By doing so, he always ends up choosing *ex post* according to his *ex ante* plans, and avoids violating dynamic consistency. In the example, at n_0 he will forecast that by taking the drug he will become an addict, and realises that his only options are *b* and *c*. Therefore he will choose the most preferred option between the two, that is, *c*.

Hammond's example of dynamically inconsistent choice allows us to consider another possible model of behaviour, which is formalised only later in the literature (McClennen 1990, Machina 1989) in the context of dynamic inconsistency under risk – *resolute choice*. The agent resolves to act according to a plan judged best from an *ex ante* perspective, and intentionally acts on that resolve when the plan imposes on him *ex post* to make a choice he does not prefer at that point. By so choosing he manages to act in a dynamically consistent manner. In Hammond's example, at n_0 a resolute agent would have resolved to act according to the plan leading to the most preferred outcome *a* – take the drug till it is harmless, then stop; and he would have acted on that resolve when at n_1 the plan imposed on him to choose the less preferred option – therefore going for *a* and not for *b*.

We have introduced the problem of dynamic inconsistencies in the context of changing preferences in a world of certainty. There are other contexts in which the same considerations apply. One is in a dynamic certain world in which agents discount the future non-exponentially – for example, using quasi-hyperbolic discounting, as in Harris and Laibson (2001). We note that Harris and Laibson remark (p 939) that "we model an individual as a sequence of autonomous temporal selves", explicitly making clear that preferences are changing through time, in the sense that the *relative* evaluation of consumption in any two periods varies depending upon when the evaluation is carried out. In this paper, and indeed in much of the literature on quasi-hyperbolic consumers, the analysis assumes that the decision maker is sophisticated in the sense that we have used it above. Hence the decision maker takes into account his or her decisions in the future when deciding in the present. In this way, the decision maker resolves his or her dynamic inconsistency.

However, it could also be the case that the decision maker acts either naively or resolutely[3] – in which cases the predictions of the quasi-hyperbolic model would be different. A further dynamic context in which exactly the same considerations apply is that of a risky world in which agents do not have Expected Utility preferences. This is the context in which we carry out our experimental investigation. We begin to give detail in the next section.

However, we have exactly the same interest in all contexts: how does a dynamically inconsistent agent respond to, and resolve, his or her dynamic inconsistencies: naively, resolutely or sophisticatedly?

2. Dynamically Inconsistent Behaviour in a Sequential Risky Context

In the context of risk and uncertainty, the problem of dynamic inconsistency is crucially linked to the question of whether the decision maker has Expected Utility preferences or not. As is well known, if the decision maker's preferences satisfy Expected Utility theory, then the problem of dynamic inconsistency does not arise. For a non-Expected-Utility decision maker, however, the problem may arise. We illustrate this in Figure 1, where three decision problems are represented in tree form. In this figure, the (green) squares represent decision nodes; and the (red) circles represent chance nodes (where Nature moves with the probabilities indicated in the figure). The letters on the various choice branches denote the implied lottery (for example K is the certainty of £30, while L is the lottery which gives £50 with probability 0.8 and £0 with probability 0.2). In Problem 2, M indicates the lottery obtained choosing Up in Problem 2 for $\varepsilon = 0$ and N the lottery for $\varepsilon = £1$. In Problem 3, M denotes the lottery obtained by playing Up at the first decision node and Down at the second while O that obtained by playing Up at both decision nodes[4].

Consider first Figure 1a and consider an individual[5] who prefers K to L in Problem 1 and O to N to M in Problem 2[6]. These preferences (for $\varepsilon = 0$) imply a violation of Expected Utility theory

[3] Or in some other way - these three ways of reacting to the potential dynamic inconsistency are not the only ones.
[4] A tree with a similar structure is found in McClennen (1990)
[5] Whom we assume here is female to avoid expositional clumsiness.
[6] Note that lottery N stochastically dominates lottery M.

in the form of a Common Ratio Effect[7]. Now consider how such an individual would tackle Problem 3, in which there are two decision nodes D_1 and D_2.

Suppose this individual is at decision node D_2. Then her preferences indicate that she would choose to move Down at that node, because she prefers K to L. Now look at the situation as viewed from node D_1. As viewed from there, if she moves Up at that first node, then she is faced with either getting O (by moving Up at D_2) or getting M (by moving Down at D_2). She prefers O, and that, by assumption, is preferred to the lottery obtained by moving Down at D_1 namely N. Hence she sets off by choosing Up at D_1. A problem arises, however, if she arrives at node D_2. As we have already argued, at this node Down is preferred, and so, if the preferences at that point are followed, the individual will choose Down at node D_2. Hence the individual, in arguing in this way, plans, at node D_1, to choose Up at node D_2 but, if she arrives there, actually chooses Down. This is dynamic inconsistency: adopting a plan but then not implementing it at a later node. Moreover, it is a problem to the individual when at node D_1. If she is aware of this dynamic inconsistency, then she realises that by choosing Up at D_1 and then choosing Down at D_2 she therefore ends up with the lottery M which is dominated by the lottery N – which she could get by choosing Down at D_1.

What might the individual do about this dynamic inconsistency? Well, she might simply be unaware of it, or ignore it, acting *naively*, and choose at each point the strategy most preferred from the perspective of *that* point - Up at D_1 and then Down at D_2. However, if she is aware of this dynamic inconsistency, she might want to do something about it. One possibility is that the individual, anticipating the fact that she will choose Down at D_2, if she reaches it, realises that it would be better to choose Down at node D_1. This, in the literature, is termed acting *sophisticatedly* – see, for example, McClennen (1990) and Machina (1989)[8]. This sophisticated individual anticipates her future behaviour, and avoids the inconsistency by making a choice at the initial decision node which is constrained by her anticipated choice at each following node.

[7] The Independence Axiom of Expected Utility theory implies that £30 for sure is preferred to a 80:20 lottery over £50 and £0 if and only if a 25:75 lottery over £30 and £0 is preferred to a 20:80 lottery over £50 and £0.

[8] Machina though does not use directly the term sophisticated choice for this kind of approach. Also Karni and Safra's model of "behavioural consistency" (Karni and Safra 1989) represents a way of implementing the sophisticated choice approach, and represents a solution to the problem of dynamic inconsistency with non-linear preferences.

However, as viewed from the individual at node D_1 this behaviour implies the lottery N whereas, as viewed from D_1, a better lottery is O which is achievable by choosing Up at D_1 and Up at D_2. The individual at D_1 might accordingly decide to be *resolute* and choose Up at D_1 and Up at D_2. This resolute individual avoids inconsistency by making the choice of a plan most preferred at the initial node to constrain future behaviour. She resolves to implement the plan originally adopted, despite the fact that this implies, at some future node, making a choice that she would not have liked to make once arrived at that node. Again a discussion of this kind of behaviour can be found in Machina (1989)[9] and McClennen (1990)[10].

Such problems do not arise with Expected Utility individuals. In the context of the above decision problem, EU preferences could be such that K is preferred to L and M to O, or the reverse. In the first case, the individual chooses Down at node D_1 and in the reverse case chooses Up at D_1 and Up at D_2[11].

3. Related Experiments

The purpose of the study reported in this paper is to try and find out whether subjects are naïve, resolute or sophisticated. To the best of our knowledge, this is the first such study. However, there are some related experiments, which also study behaviour in dynamic contexts. We shall confine ourselves to those in economics which have an appropriate incentive mechanism[12]. One such paper is that by Cubitt *et al* (1998) which is primarily concerned with trying to discover, in the context of dynamic decision making, the axiom of Expected Utility theory which is apparently violated by individuals with non-EU preferences. Cubitt *et al* presented (to different groups of subjects) different decision trees and studied the behaviour of subjects in such trees. These trees were similar to those in Problems 1, 2 and 3 shown in Figure 1. A similar set of trees was used in

[9] Machina's model of choice is equivalent even if differently formalised to McClennen's model of resolute choice. According to Machina, resolute choice represents one of the "antecedents of the formal approach" presented in his paper.

[10] An interesting pratical point is how this person can force him or her-self to behave resolutely.

[11] We should, for completeness, also mention the case when L is preferred to K and N to O to M. In this case the individual chooses Down at D_1 and no dynamic inconsistencies arise.

[12] Busemeyer and his associates in psychology have carried out some related experiments (see, for example, Busemeyer *et al* 2000) but without such incentives.

Hey and Paradiso (2006) who, instead of looking at the *decisions* of subjects, collected data on subjects' *evaluations* of decision problems. The two studies are closely related though they had different objectives. These three sets of trees are relevant to the objectives of this present paper since naïve, resolute and sophisticated subjects could have different behaviours in the trees, as well as having different evaluations of them. Cubitt *et al.* found that subjects behave differently in the different trees, while Hey and Paradiso find that the temporal frame also affects the evaluations the subjects make of the different trees.

Our experiment differs from both these previous studies, most particularly because this present paper has a different agenda. However, we build on the framework set by these two papers by using similar decision trees to those used in Cubitt *et al* and in Hey and Paradiso. We also add a fourth tree which helps us in our task of distinguishing between naïve, resolute and sophisticated subjects. In addition, our experiment differs from both these other two in that we observe not only behaviour but also gather data on valuations. In a sense, our experiment is a 'cleaned-up' combination of these other two experiments but with a differently directed research agenda – that of detecting whether subjects are naïve, resolute or sophisticated.

Our experiment is built around the four trees shown in Figure 2. It will be noticed that Tree 1 is identical to Problem 2 of Figure 1a; Tree 2 is effectively the same as Problem 1 in Figure 1a – with the addition of some preliminary risk; Tree 3 is the same as Problem 3 as shown in Figure 1b; and Tree 4 is a tidied-up version of Tree 3, in which the decision problem is reduced to that of a single decision – rather than that of a sequence of decisions. We note that Expected Utility theory has a strong prediction about relative behaviour in Trees 1 and 2; given that, as we have discussed above, we have predictions about the behaviour in Tree 3 of naïve, resolute or sophisticated subjects whose preferences do not respect Expected Utility theory.

In essence, observing behaviour in these trees should enable us to detect whether people are naïve, resolute or sophisticated. However, there are two problems. The first is that we can only hope to detect differences in type if the agent has non-EU preferences – since the distinction only has sense for non-EU agents. Moreover, it may be the case that a subject has non-EU preferences but our trees do not enable us to detect this fact: a subject who chooses K in Problem 1 and M in Problem 2 (or L in Problem 1 and O in Problem 2) is not necessarily a person with EU preferences. A second problem is that, inevitably in experiments, there is *noise* in the subjects' responses. It is well known from many previous experiments that this noise can be quite large; to give some idea of this, we note that, when asked the same question on two separate occasions, subjects give different answers roughly 25-35% of the time. Accordingly we cannot be sure that any stated decision (or any stated valuation) is *exactly* in accordance with the subject's preferences.

To help us to get over these two problems, we used *three* sets of four trees – all with the same structure as the four trees in Figure 2 but with different probabilities and payoffs. We give detail in the next section (section 4). Moreover, in analysing the data from the experiment, we explicitly assumed the existence of noise in the responses of the subjects and we use all the data on each subject to analyse the preferences of that subject. We give detail in section 5.

4. The experimental design and implementation

The experiment was built on the four decision problems introduced in the previous section. We used three sets each with four trees, with different values for the payoffs and the probabilities, giving a total of 12 decision trees. Table 1 gives the details; we note that the possible payoff varies from a minimum of £0 to a maximum of £150. We designed the experiment in such a way that we could observe subjects' decisions in all 12 trees as well as their stated evaluations of the 12 trees. The amounts of money and the probabilities in the three sets of trees were chosen to satisfy various criteria. First, we wanted one set of trees (Set 1) that was the same as used in previous experiments (Cubitt *et al* 1998 and Hey and Paradiso 2006). Second, we wanted different amounts of money, and different probabilities in the different sets, but with the same properties. Hence the structure of the payoffs and probabilities are identical in the three sets – they are always such that we should be able to detect non-EU behaviour from the choices, and then to distinguish the naïve, resolute and sophisticated types. Finally we wanted payoffs and probabilities that gave an incentive to the subjects that was roughly similar in the three different sets (the maximum expected payoff ranges from £7.50 to £10 in the different sets) as well as being reasonable. Here there is a problem in that we preferred that subjects would not leave the experiment with less money than that with which they came. Obviously there is no way to guarantee this, but giving a relatively high participation fee reduces this possibility – and also increases the amount of money that they can bid for the various trees. For a risk-neutral subject, the value of a randomly chosen tree in the different sets is £8.83, so while the value of bidding for a tree was not high, it was also not negligible. Moreover subjects were encouraged to think seriously about their bids by having to spend 15 minutes deciding on them, which turned out to be a more than sufficient length of time.

The experiment was conducted at **EXEC**, the Centre for Experimental Economics at the University of York. A total of 50 students, both graduate and undergraduate, took part in the experiment. They were given written instructions (see Technical Appendix 4). When all participants had finished reading the instructions, a PowerPoint presentation was played at a predetermined speed on their individual screens. After this, they could ask questions. The experiment then started, using a Visual Basic program (available on request). The computer screen showing the four decision trees is that shown in Figure 2. In order to elicit the subjects' evaluations for each of the twelve trees, we used the second-price sealed-bid auction method as in Hey and Paradiso (2006). This was implemented as follows. Subjects performed the experiment in groups of five. They were sat at individual computer terminals and were not allowed to communicate with each other. They individually made bids for each of the three sets of four trees (twelve trees overall) and were given 15 minutes to bid for each set of four trees in the set. During the bidding period the subjects were allowed to practice playing out the decision trees as much as they wanted in the time allowed. It was made clear that the outcomes of the practice did not affect their payments in any way. The number of seconds left for the practice and bidding was shown in the box at the top left-hand side of each decision tree. When the bidding time was over, the subjects played out all the twelve problems for real. We displayed on each subject's screen the results of his or her playing out plus the bids of all the 5 subjects in the group for each of the twelve trees. Then we invited one of the 5 subjects to select a ball at random from a bag containing 12 balls numbered from 1 to 12. This determined the problem on which the auction was held. We then consulted the bids made by the group members, and the subject with the highest bid for the problem paid us the bid of the second highest bidder. As all subjects were given a £20 participation fee, 4 of the 5 members earned £20, while the fifth earned £20 minus the bid of the second highest bidder plus the outcome. Both in the instructions and in the presentation, it had been emphasised, through different examples, that the highest bidder (in the auction corresponding to the randomly selected decision problem) would have to pay the bid of the second highest bidder, and it was made clear that the bid for each of the twelve problems should be equal to the willingness to pay for the decision problem. It was emphasised that in the case that the subject's bids were all £0, he or she was unlikely to be sold one of the decision problems and thus will definitely end up with no less than the participation fee. Moreover, in the case that the bid was higher than £0 and the subject's bid was the highest for the chosen tree, he or she would end up with the participation fee, minus the bid of the second highest bidder, plus the outcome. This could – depending on the bid – be less than the participation fee and the subject could end up losing money. Subjects were warned in advance that they had to bring enough cash to the experiment to allow for this possibility.

5. Analysis of the data from the experiment

If there was no noise in subjects' behaviour, the analysis of the data would be simple. First we would select only those subjects for whom the decisions in Trees 1 and 2 were inconsistent with EU theory. We would then see what decisions they took in Tree 3; this would reveal (for those subjects who chose Down in Tree 1 and Up in Tree 2) whether their behaviour was consistent with one or other of the three specifications: naïve, resolute or sophisticated. For example, if they chose Down in Tree 1 and Up in Tree 2 (thus revealing their inconsistency with EU), then their behaviour in Tree 3 would reveal their type: naïve if they moved Up and then Down, resolute if they moved Up and then Up; and sophisticated if they moved Down.

However, there is noise in the data. This creates an immediate problem in that their behaviour might suggest that they were naïve on Set 1, resolute on Set 2 and sophisticated on Set 3[13]. What would we conclude then? More seriously, because of their noise, we cannot be sure whether behaviour in Trees 1 and 2 actually does reveal EU or non-EU preferences. They may be EU but choose Up in Tree 1 and Down in Tree 2 because of this error; similarly they may be non-EU and yet choose Up in both Trees 1 and 2 (either with or without error). We need to take into account of this noise. What we do then, rather than carry out a series of individual tests on the data which require the assumption of zero noise, we use all the data to estimate models of behaviour. Doing so enables us to use all the data on each subject (24 observations), and, in particular, to try and discover the type (naïve, resolute or sophisticated) of each subject. Methodologically, it seems better to use all the data on each subject to try and fit the various types rather than to carry out a series of individual tests on bits of the data.

Our data analysis strategy is the following. We assume that subjects are different, both in terms of their preferences and in terms of their *type* (naïve, resolute or sophisticated). We fit specifications to the data from each subject individually. We use all the data (24 observations) on each subject with our ultimate goal of telling whether the subject is (more likely[14] to be) naïve, resolute or sophisticated.

While the analysis of the data on decisions only requires us to know preferences over a relatively small number of lotteries, the analysis of the data on evaluations of the trees requires us to know preferences not only over these lotteries but also over all the evaluations of these lotteries (in monetary terms by the subject). For each subject, there are 12 of these evaluations (the bids made for each of the twelve trees). In each set of four trees, there are 10 lotteries, involving 5 amounts of money. We neither know which type the subject is, nor the preference functional of the subject.

[13] Recall that subjects were presented with three sets of trees with the structure as in Figure 2.
[14] We note that the presence of this error means that inevitably our inferences can not be certain.

This has to be estimated. Given that we have just 24 observations per subject, it is clear that we cannot estimate the evaluation of each of the 30 lotteries (10 lotteries in each set of trees) involved in the experiment; we have to make some restrictions. Accordingly we assume a particular form of the preference functional – to be precise that of the Rank Dependent Expected Utility model. This seems to be well accepted in the literature as the empirically most-valid generalisation of EU; moreover, it contains EU as a special case.

The Rank Dependent functional is composed of a utility function and a probability weighting function; we denote the former by *u(.)* and the latter by *w(.)*. The Rank Dependent Expected Utility, *U(G)*, of a gamble $G = (x_1, x_2, \dots, x_I; p_1, p_2, \dots, p_I)$, where the prospects are indexed in order from the worst x_1 to the best x_I, is given by

$$U(G) = u(x_1) + \sum_{i=2}^{I} [u(x_i) - u(x_{i-1})] w(p_i + p_{i+1} + \dots + p_I) \tag{1}$$

We note that Rank Dependent Expected Utility preferences reduce to Expected Utility preferences when the weighting function is given by $w(p) = p$.

To fully characterise the preferences of a subject obeying the Rank Dependent Expected Utility model, we need to know the utility function *u(.)* and the weighting function *w(.)*. As we describe later, we assume particular functional forms for these two functions and estimate the parameters of the functions. We choose the functional forms which best fit the responses of each subject – in a way that we will describe shortly.

We now need to have a story about the noise in the subjects' responses – more technically, we need to specify the stochastic structure of the data. There are various stories that one can use and we choose to follow a Fechnerian measurement error story[15]. To be precise, we assume that, when

[15] Before arriving at this particular specification, we tried several others, of varying degrees of sophistication. One simple alternative was that subjects made all evaluations with error but then 'trembled' (see Moffatt and Peters 2001) when taking decisions and when making bids; the trouble with this story (in addition to the fact that it does not seem empirically valid) is that, while the tremble story is simple to apply to decisions, it is not obvious how to interpret it with respect to bids. There are also other variations that we have tried on the basic story that we report in this paper; in particular, we explored the hypothesis that subjects made no mistakes when evaluating certainties - this performed worse than the variant reported in the paper; and also a variant that takes into account that the extreme value distribution incorporates a bias (the expected value of a variable with an extreme value distribution with parameters *m* and *1/s* is not

evaluating any lottery (whether certain or risky), the subject makes a measurement error. More specifically, if *u(.)* is the utility function of the individual, and $u^{-1}(.)$ its inverse, then we assume that the evaluated certainty equivalent of any gamble *G* is given by $u^{-1}(U(G)) + e$, where *U(G)* is the Rank Dependent Expected Utility of the gamble (using equation (1)) and where *e* is an error – a measurement error. To complete the story, we need to specify the distribution of *e*. We do this in a way that is tractable and not unreasonable – we assume that *e* has an Extreme Value distribution with parameters *0* and *1/s*. Thus the cumulative distribution function, *F(.)* of *e* is given by:

$$F(e) = \exp(-\exp(-es)) \tag{2}$$

It follows that the probability density function *f(.)* is given by:

$$f(e) = \exp(-\exp(-es))\exp(-es)s \tag{3}$$

To summarise: we assume that the subjects each have a well-defined (Rank-Dependent) preference function and that each subject is either naïve, resolute or sophisticated. For each subject we find the best-fitting preference function and the best-fitting specification (naïve, resolute or sophisticated) to represent the subject's responses on the experiment. The responses are the bids (for the 4 trees in each of the 3 sets) and the decisions (in 4 trees on each of the 3 sets) – a total of 24 observations[16] for each subject. We note that the nature of the data is different – for the bids we have a number, while for the decisions we have their choice. The former is essentially a continuous variable while the latter is a discrete variable (taking 1 of either 2 or 3 values). The analysis of the two kinds of data has to be different. The details are given in the next section, where we are more specific about how we have interpreted and applied our stochastic specification.

6. The Stochastic Specification

We have assumed that subjects evaluate lotteries with error. More specifically we have assumed that the expressed valuation, *VG*, of a lottery *G* with true certainty equivalent $u^{-1}(U(G))$ is

---

m but rather $m+\gamma s$ where γ is the Euler-Mascheroni constant 0.5772156649), by exploring the notion that the subjects corrected for this bias when making their bids; this also performed worse than the variant we have used in the paper.

[16] We note that we deliberately fixed the random number generator in the real playing out so that we could have observations on their choices later in the tree. We did not, however, fix the generator during the practice playing out. We note to our shame that this is deception, but deception that is in the subjects' interests.

given by $VG = u^{-1}(U(G)) + e$ where e has an extreme value distribution with parameters 0 and $1/s$. When taking a decision between lottery A or lottery B, we assume that the subject evaluates both A and B with error and chooses the lottery with the highest evaluation. Given the properties of the extreme value distribution, it follows that the probability that A is chosen is $\frac{e^{VAs}}{e^{VAs}+e^{VBs}}$ and the probability that B is chosen is $\frac{e^{VBs}}{e^{VAs}+e^{VBs}}$ where VA is the true valuation of lottery A and VB the true valuation of lottery B. We note the importance of the parameter s: when s is zero then both lotteries are chosen with probability one-half, and when s is infinite then the lottery with the highest true valuation is chosen with certainty. We can therefore interpret s as indicating the precision of the subject's decision – the higher is s the more precise is the subject, and the more likely he or she is to choose the lottery with the truly highest value.

The above story can easily be extended to a choice between three lotteries A, B and C: given our error story, lottery Z (= A, B or C) is chosen with probability $\frac{e^{VZs}}{e^{VAs}+e^{VBs}+e^{VCs}}$.

This error story also enables us to determine the probability density of the bids. We assume that if a tree offers two choices – between lotteries A and B – then the bid is put equal to the maximum of the evaluations of A and B. Using the properties of the extreme value distribution, and assuming that the errors in the valuations are independent, it follows that the cumulative distribution function of the bid x is given by

$$\exp\{-[\exp(-(x-VA)s)+\exp(-(x-VB)s)]\} \tag{4}$$

From equation (4) we can find the probability density of the bid x based on two lotteries.

An obvious extension leads us to the cumulative distribution function of the bid when the tree offers three lotteries A, B and C and the bid is put equal to the maximum of the evaluations of these three lotteries:

$$\exp\{-[\exp(-(x-VA)s)+\exp(-(x-VB)s)+\exp(-(x-VC)s)]\} \tag{5}$$

From this we can find the probability density of the bid x based on three lotteries.

We now need to say something about how the different types are assumed to behave in the four trees. Tree 1 is simple and all types do the same thing. The decisions are based on the evaluations of M (Up) and O (Down), and the bids are made on the basis of the maximum of these two evaluations.

In Tree 2 different types do different things. The naïve subject takes the decision at the decision node and this decision is based on the evaluations of K (Up) and L (Down); however the bid is based on the evaluations as viewed from the beginning of the tree – from which perspective the choice is between M and O – and hence the bid is based on the maximum of the evaluations of M and O. The resolute subject, however, bases both the decision and the bid as viewed from the beginning of the tree – that is on the basis of the evaluations of M and O. The sophisticated works backwards: he or she anticipates that the decision at the decision node will be based on the relative evaluations of K and L; if K is chosen he evaluates the tree as being worth the evaluation of M; if L is chosen the tree is worth the evaluation of O. He or she can work out the probability of choosing K and L and hence work out the expected evaluation of the tree – on which he or she bases his bid.

In Tree 3 again the different types do different things. The actual decision of a naïve subject at the second decision node is based on the evaluations of L and K, but the naïve agent does not anticipate this when taking the decision at the first node and in making the bid for the tree. These are determined by the evaluations of M, N and O. If either the evaluations of M or O are bigger than the evaluation of N the naïve agent chooses Up at the first node; otherwise he or she chooses Down. The bid of a naïve is simply determined by the maximum of the evaluations of M, N and O. In contrast a resolute subject bases both his decisions at both nodes on the evaluations of M, N and O: if the evaluation of M is the highest he or she chooses Up and then Down; if the evaluation of O is the highest he or she chooses Up and then Up; and if the evaluation of N is the highest, he or she chooses Down; the bid of a resolute subject is based on the maximum of the evaluations of M, N and O. The sophisticated subject once again anticipates his or her decision at the second node – this will be Up if L is evaluated more highly than K and Down otherwise. If L is evaluated more highly than K then his or her decision at the first node (and the bid for the tree) is based on the evaluations of O and N; if, however, K is evaluated more highly than L then his or her decision at the first node (and the bid for the tree) is based on the evaluations of M and N.

Tree 4 is simple and all do the same: the decisions and the bids are all based on the evaluations of M, N and O.

There is one slight complication that we have so far ignored: and that concerns what exactly a sophisticated person does when working backwards – when backwardly inducting. We have assumed so far that this agent eliminates the branches of the tree that he or she will not be following in the future, and then simplifies the tree by using *reduction* – that is, reducing the remaining compound lottery to a simple lottery using the usual probability rules. But there is an alternative – that this agent simplifies the tree by substituting in *certainty equivalents* of the remaining bits of the tree. Segal (1999) notes and discusses this distinction. Accordingly we should distinguish between two types of sophisticated agents – those who work backwards using reduction and those who work backwards using certainty equivalents. We refer to these as Type 1 sophisticated and Type 2 sophisticated respectively. This distinction does not affect the story we have told above about how sophisticated agents process a dynamic tree, but it does affect the specification of the appropriate likelihood functions.

We confine all the technical details to Technical Appendix 1 (with the GAUSS program for one of the combinations in Technical Appendix 2). Suffice it to say here that we estimate the parameter of the utility function, the parameter of the weighting function and the precision parameter s using maximum likelihood (implemented in GAUSS), subject by subject and specification by specification. To do this we need to specify the likelihood of the observations. This is different for the decision and for the bids: in essence the likelihood of a decision is the probability that that decision is taken given the parameters and the stochastic specification; the likelihood of a bid is the probability density at the bid given the parameters and the stochastic specification.

There is one further problem that we need to consider – the specifications of the utility function and of the weighting function. Ideally one would not specify functional forms but estimate the functions at all possible values. The problem with this is, as we have already noted, that there are too many values to make this procedure possible: given the number of observations that we have (24 for each subject), we would lose too many degrees of freedom. So we are forced to assume functional forms. The most obvious ones are the CARA and CRRA specifications for the utility function ($u(x) \propto -\exp(-rx)$ and $u(x) \propto x^r$ respectively) and the Quiggin (1982) and Power specification for the weighting function ($w(p) = \frac{p^g}{(p^g + (1-p)^g)^{1/g}}$ and $w(p) = p^g$ respectively). The Quiggin specification allows for an S-shaped weighting function while the Power specification does not. In order to ensure the robustness of our results we use all four possible combinations: CRRA with Power; CRRA with Quiggin; CARA with Power; and CARA with Quiggin, and for each subject we choose the combination which gives the highest maximised log-likelihood averaged over all four (naïve, resolute Type 1 and Type 2 sophisticated) specifications[17].

7. The Results

To summarise: we have for each subject fitted the naïve, resolute and (the two types of) sophisticated specifications to the 24 observations for each subject for each of the four combinations of utility and weighting function. The detailed estimates are available on request. An illustration of some of the results for one of the four combinations (CARA with Quiggin) is provided in Technical Appendix 3. For each type and for each specification we have the following information from our estimations:

- The value of the maximised log-likelihood;
- Information as to whether the maximum likelihood converged correctly;

[17] We should note that for 23 subjects there was one combination which fitted best on all four specifications; there were 24 subjects for whom one combination fitted best on three of the four specifications; and there were 3 subjects for whom one combination fitted best on two of the four specifications. The conclusion seems to be clear - for virtually all subjects, the data seems to be telling us that one combination (of utility function and weighting function) fits best *independently* of the specification.

- The estimates and standard errors of (the transformed values[18] of) the parameters - *r, g* and *s*.
- (Hence) The estimates of the parameters.

We should note that for some of the specifications there were occasional problems with the convergence of the maximum likelihood process, but there were combinations (particularly CARA with Quiggin) that seemed to be particularly robust. We should note that there is no guarantee that the likelihood function is smoothly concave everywhere; accordingly we tried several starting values for the maximum likelihood software. We are reasonably convinced that we have found the overall maximum in essentially all cases.

The bottom line is the categorisation of subjects as to whether they are naïve, resolute or sophisticated. As we have already noted, we begin by choosing, for each subject, the combination of utility function and weighting function that best explains the data on the subject. In some cases this was simple: irrespective of the type (naïve, resolute or sophisticated) the same combination yielded the maximum of the maximised log-likelihoods - there were 25 subjects for whom the data had this property. For the rest we had to select the best combination in some manner; our criterion was to select the combination for which the average maximised log-likelihood (across all types) was maximised. This may seem somewhat arbitrary but we have carried out an independent check by seeing if the chosen type might have been different if we had chosen an alternative (but possible) combination. We found that, for all except 9 of the 50 subjects, the maximised log-likelihood with our chosen 'best' combination and chosen 'best' type was, in fact, the highest of all the 16 maximised log-likelihoods that we computed; for 7 of these 9, the maximised log-likelihood of our chosen 'best' type was higher than for any other type in all the other combinations. Just in two cases (subjects 11 and 30) would our chosen combination change our chosen type – subject 11, whom we have classified as more likely to be naïve, could be classified as being sophisticated; and subject 30,

[18] To avoid problems with the maximum likelihood algorithm we transformed our parameters to restrain their range (for example to stop the *s* parameter becoming negative). The returned estimates and variance-covariance matrix are thus those of the transformed parameters - and they need to be transformed back before they can be interpreted.

whom we have also classified as being naïve, could possibly be reclassified as resolute. We feel that our classification of our types is reasonably robust.

We give some examples of what we have done. Consider Table 2 in which we give examples of three subjects which show different characteristics. All the entries in the tables are maximised log-likelihoods. The rows indicate the *type* (naïve, resolute or the two types of sophisticated) and the columns indicate the *combination* (of utility function and weighting function) used in the estimation.

Table 2a shows the log-likelihoods for a subject for whom the best type is the same irrespective of the combination. (As we have already noted 25 of the 50 subjects had this property.) Indeed, for this subject the CARA/Power combination performs better than the CARA/Quiggin combination and those two significantly better than either of the combinations involving the CRRA utility function. Moreover, independently of the combination, the naïve specification fits the data better than the others. We therefore select the CARA/Power combination as representing the subject's preferences and conclude that this subject is more likely to be naïve than any of the other types.

Subject 3, shown in Table 2b, is somewhat less clear cut. For this subject the best fitting combination depends upon the type; only by averaging across the types are we able to declare the CARA/Quiggin combination 'best' for this subject. There were 20 subjects for whom the estimates had this property. However, we do note that, whatever the combination, the log-likelihood for the resolute type is higher than for the other types. We therefore conclude that this subject is resolute.

Subject 46, shown in Table 2c, is even less clear cut. Only 5 of the 50 subjects were similar to this one. It will be seen that for two of the types (naïve and resolute) the CRRA/Power combination is best, whereas for the two sophisticated types the CARA/Power combination is best. However, both sophisticated types are best for three of the combinations and it is only for the CRRA/Power combination does the sophisticated type come out worse than the other two types. To decide on the best type for this subject, we have averaged the log-likelihoods over all types for each combination and have chosen the combination for which this average is highest – this leads us to the choice of the CARA/Power combination as best representing this subject's preferences. Given that, we note that the sophisticated likelihoods are the largest in the CARA/Power combination – and hence we declare 'sophisticated' as the winner. Note too, that of all 16 maximised log-likelihoods in the table (excluding the means) those for the sophisticated with the CARA/Power combination are the largest.

At this point we have chosen the best combination for each subject. At the same time we have also implicitly chosen the best type for each subject (that is, the type for which the maximised log-likelihood for the chosen combination is greatest). However we also want to give an indication of how much better the best-fitting type is relative to the others. To this end we proceeded as follows. For each subject we have three maximised log-likelihoods (where we take just one of the two Types of sophisticated specifications - that with the highest maximised log-likelihood). Let us denote these by *ll(n), ll(r)* and *ll(s)* for the naïve, resolute and sophisticated specifications respectively. If we adopt a Bayesian interpretation of the results, and if we start with equal priors on the three specifications, then the posterior probabilities of the naïve, resolute and sophisticated specifications being the correct ones are respectively

$$P(i) = \frac{\exp(ll(i))}{\exp(ll(n)) + \exp(ll(r)) + \exp(ll(s))} \qquad i = n, r, s \tag{6}$$

Hence, for example, for Subject 1[19] on the CRRA plus Quiggin combination we have:

ll(n) = -11.78101, ll(r) = -14.15611 and *ll(s) = -12.49711*

and hence, applying formula (6), we have that the posterior probabilities for the three specifications are

P(n) = 0.632, P(r) = 0.059 and *P(s) = 0.309.*

Hence, for this subject the naïve specification is over twice as likely as the sophisticated specification, and the resolute specification is relatively very unlikely.

[19] In this case the maximised log-likelihoods for the Type 1 and Type 2 sophisticated specifications were -13.32295 and -12.49711 and we accordingly prefer Type 2 to represent sophistication.

We have applied this analysis to each of subjects (using the best fitting combination of utility function and weighting function for each subject individually) and present the results graphically in Figure 3. In this triangle we represent the probability of the naive specification being correct on the horizontal axis, and the probability of the sophisticated specification being correct on the vertical axis. The probability of the resolute specification being correct is the residual. In each of these triangles subjects are indicated by a number. The triangles are divided into three areas – the one to the top being where the sophisticated specification is most probable, the one to the right being where the naive specification is most probable and the one nearest the origin being where the resolute is most probable. We note that there are 25 subjects in the "naive most likely area", 20 subjects in the "resolute most likely area" and just 5 subjects in the "sophisticated most likely" area.

Of course, for subjects whose preferences are those of Expected Utility, there is no possibility of dynamic inconsistency and hence no meaning to the distinction between the different specifications (naïve, resolute and sophisticated). In principle, therefore, we should exclude such subjects from our analysis. However, given the stochastic nature of our data we cannot be sure whether a subject is EU or not, though we do have estimates of the parameter g of the weighting function. If g is equal to 1, the weighting function reduces to $w(p) = p$ and the Rank Dependent model reduces to Expected Utility. We can carry out formal tests (based on our estimates) of the hypothesis that the g parameter is significantly different from 1. Out of the 50 subjects there are 28 subjects for whom the g parameter is significantly different from 1 at the 1% level - and hence probably not subjects with Expected Utility preferences. If we restrict our analysis to these 28 subjects we get Figure 4. It will be seen that there are 15 in the "naive most likely" area, 12 in the "resolute most likely" and just 1 in the "sophisticated most likely". These results strengthen our comments above.

We conclude that the sophisticated specification performs consistently worse than the other two, and that the naïve specification performs marginally better than the resolute specification. We comment further on these findings in the next and concluding section.

8. Conclusions

Dynamically inconsistent economic agents provide a serious challenge to economic theory. In order to predict their future behaviour, one needs to know how they are resolving any dynamic decision problem – in particular, for example, whether they are naïve, resolute or sophisticated. In conducting experiments to try to determine whether subjects are naïve, resolute or sophisticated, experimenters have a problem in designing the experiment in such a way that they can infer what subjects are doing when processing the dynamic decision problem. Essentially, the issue is concerned with whether subjects are planning their future behaviour and anticipating any possible future inconsistencies. Ideally, one would like to design an experiment where any plans that the subject is making are revealed. However, there are serious (possibly insurmountable[20]) difficulties in observing plans and hence in seeing whether they are implemented. If one simply asks the subject what he or she is planning to do, then, unless one forces the subject to implement that plan, there is no incentive for the subject to report any plans honestly; moreover, asking them such a question raises in the subject's mind the idea that they possibly *ought* to plan. Furthermore, if one insists that the subject implements the plan that he or she has stated, then one changes the nature of the problem from a dynamic problem to a pre-commitment problem.

The experiment surmounts these difficulties with a unique design in which not only behaviour is observed but also reported evaluations of different dynamic decision problems are obtained. Moreover, we have constructed the decision problems in such a way that we can distinguish, using data on both decisions and reported evaluations, the decision process actually followed.

In particular, we can distinguish between naïve decision makers (those who ignore any possible future inconsistencies), resolute decision makers (those who are resolute in implementing their *a priori* plans) and sophisticated decision makers (who anticipate their future inconsistencies

[20] But see Busemeyer *et al* (2000) for the psychologists' way round the problem.

and who are not sufficiently resolute to overcome them). These different types will have different behaviour in our trees as well as different evaluations of the trees. We have used the data to infer the type of each subject. The picture is somewhat clouded as we do not know *ex ante* the preference functional of each individual, but we have investigated four different combinations for each subject and chosen the best. While the final picture is not totally clear, it seems to be the case that around 50% of our 50 subjects are naive, 40% are resolute and just 10% sophisticated.

The large number of resolute subjects and the small number of sophisticated subjects in our experiment surprised us, as we thought *ex ante* that it would be difficult for subjects to be resolute. However, taking into account that we obliged the subjects to spend 15 minutes evaluating the trees (and practising playing them out) it seems to have been the case that subjects used this time to work out an *ex ante* strategy and to realise that it was better for them to implement that rather than behaving in a sophisticated way: remember that sophisticated behaviour is not optimal as viewed by the decision maker at the beginning of the tree. It would be interesting to explore whether this finding is sensitive to the length of the time spent contemplating the trees.

The implications for economic theory are significant. If we look at models which incorporate dynamically inconsistent behaviour (such as the literature on quasi-hyperbolic discounting in the context of a life-cycle saving model[21]), it will be seen that most of these models assume sophisticated behaviour. Our results suggest that this might be descriptively implausible. If subjects are indeed naïve or resolute rather than sophisticated, then the predictions of these models need to be modified appropriately.

[21] See, for example. Harris and Laibson (2001).

References

Beggs, S., Cardell, S. and Hausman, J. (1981), “Assessing the Potential Demand for Electric Cars, *Journal of Econometrics*, 16, 1-19.

Busemeyer, J.R., Weg, E., Barkan, R., Li, X., and Ma, Z. (2000), “Dynamic and consequential consistency of choices between paths of decision trees”, *Journal of Experimental Psychology: General*, 129, 530-545.

Cubitt, R.P., Starmer, C. and Sugden R. (1998), “Dynamic Choice and the Common Ratio Effect: An Experimental Investigation”, *Economic Journal*, 108, 1362-1380.

Cubitt, R.P., Starmer, C. and Sugden, R. (2004), “Dynamic Decisions Under Uncertainty: Some Recent Evidence from Economics and Psychology,” in *The Psychology of Economic Decisions, Volume 2: Reasons and Choices* edited by Brocas, Isabelle and Carrillo, Juan D., Oxford University Press.

Hammond, P.J. (1976), “Changing Tastes and Coherent Dynamic Choice”, *Review of Economic Studies*, 43, 159-173.

Hammond, P.J. (1977), “Dynamic Restrictions on Metastatic Choice”, *Economica*, 44, 337-350.

Hammond, P.J. (1988a), “Consequentialism and the Independence Axiom”, in Munier, B.R. (ed), *Risk, Decision and Rationality*, Dordrecht, Reidel, pp. 503-516.

Hammond, P.J. (1988b), “Consequentialist Foundations for Expected Utility”, *Theory and Decision*, 25, 25-78.

Hammond, P.J. (1989), “Consistent Plans, Consequentialism, and Expected Utility”, *Econometrica*, 57, 1445-1449.

Harris, C. and Laibson D. (2001) “Dynamic Choices of Hyperbolic Consumers”, *Econometrica,* 69, 935-957.

Hey, D.J. and Paradiso, M. (2006), “Preferences Over Temporal Frames in Dynamic Decision Problems: An Experimental Investigation”, *The Manchester School*, 74, 123-137.

Karni, E., Safra, Z (1989), "Dynamic Consistency, Revelations in Auctions, and the Structure of Preferences", *Review of Economic Studies*, 56, 421-433.

Machina, M.J. (1989), "Dynamic Consistency and Non-Expected Utility Models of Choice Under Uncertainty", *Journal of Economic Literature*, 27, 1622-1668.

McClennen, E.F. (1990), *Rationality and Dynamic Choice: Foundational Explorations*, Cambridge: Cambridge University Press.

Moffatt, P.G. and Peters, S.A. (2001), "Testing for the presence of a tremble in economic experiments", *Experimental Economics*, 4, 221-228.

Rabin, M. (1998), "Psychology and Economics", *Journal of Economic Literature,* 36, 11-46.

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

Raiffa, H. (1968), *Decision Analysis*, Addison-Wesley, Reading, MA.

Sarin, R. and Wakker, P. (1998), "Dynamic Choice and Non Expected Utility", *Journal of Risk and Uncertainty*, 17, 87-119.

Segal, U. (1999), "Dynamic Consistency and Non-Expected Utility", in L. Luini (ed), *Uncertain Decisions. Bridging theory and experiments*, Dordrecht: Kluwer.

Strotz, R. H. (1955-6), "Myopia and Inconsistency in Dynamic Utility Maximisation", *Review of Economic Studies*, 23, 165-180.

Table 1: The parameters used in the experiment

Parameter	Set 1	Set 2	Set 3
a	£50	£150	£60
b	£31	£51	£41
c	£30	£50	£40
d	£1	£1	£1
e	£0	£0	£0
q	0.8	0.5	0.75
r	0.25	0.1	0.20

Table 2: Examples of log-likelihoods

Table 2a: Subject for whom best combination is the same irrespective of the type

Subject number 8	aq	ap	rq	rp	best	highest log-likelihood
n	-1.85533	-0.50035	-13.95100	-15.35401	ap	-0.50035
r	-1.98967	-0.54727	-18.18127	-18.18191	ap	-0.54727
s1	-1.98032	-0.54233	-14.72916	-15.66046	ap	-0.54233
s2	-8.73564	-0.54233	-14.71678	-15.66046	ap	-0.54233
mean	-3.64024	-0.53307	-15.39455	-16.21421	ap	-0.53307

Winner is naïve.

Table 2b: Subject for whom best combination is the same for three of the types

Subject number 3	aq	ap	rq	rp	best	highest log-likelihood
n	-15.62822	-16.13782	-15.88433	-17.09272	aq	-15.62822
r	-14.03244	-13.21857	-14.03246	-13.68696	ap	-13.21857
s1	-15.69680	-16.00562	-15.85117	-16.59095	aq	-15.69680
s2	-15.75262	-16.00562	-15.94487	-16.59095	aq	-15.75262
mean	-15.27752	-15.34191	-15.42821	-15.99040	aq	-15.27752

Winner is resolute

Table 2c: Subject for whom best combination is the same for two of the types

Subject number 46	aq	ap	rq	rp	best	highest log-likelihood
n	-21.64725	-21.13685	-22.07941	-20.74374	rp	-20.74374
r	-22.56506	-21.83636	-22.96630	-20.81319	rp	-20.81319
s1	-21.28023	-20.36172	-22.01664	-21.27399	ap	-20.36172
s2	-20.82796	-20.36172	-21.37605	-21.27399	ap	-20.36172
mean	-21.58013	-20.92416	-22.10960	-21.02623	ap	-20.92416

Winner is sophisticated (either Type)

Glossary

Types:

n naïve
r resolute
s1 Type 1 sophisticated
` s2 Type 2 sophisticate

Combinations:

aq CARA and Quiggin
ap CARA and Power
rq CRRA and Quiggin
rp CRRA and Power

Figure 1: The problem of dynamic inconsistency

Figure 1a: The violation of expected utility preferences

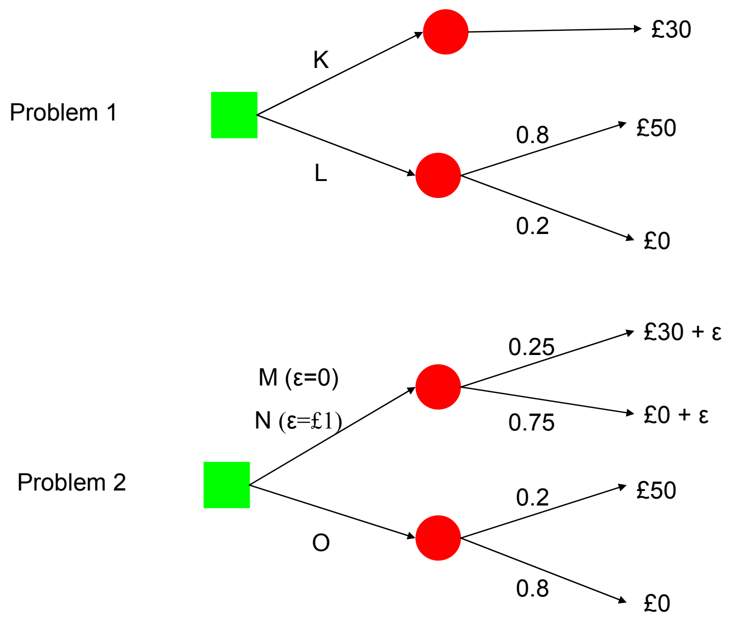

Figure 1b: The implications for dynamic choice

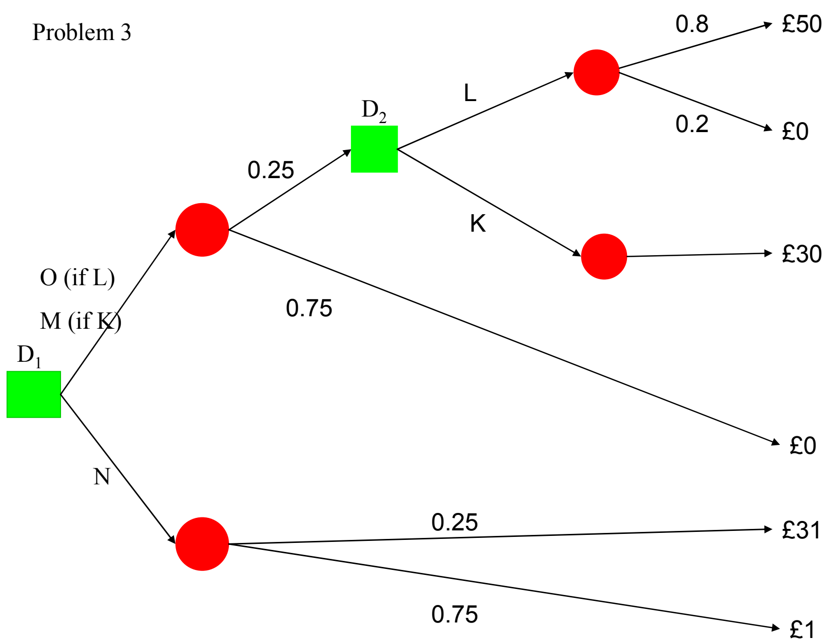

Figure 2: A screen shot from the experiment

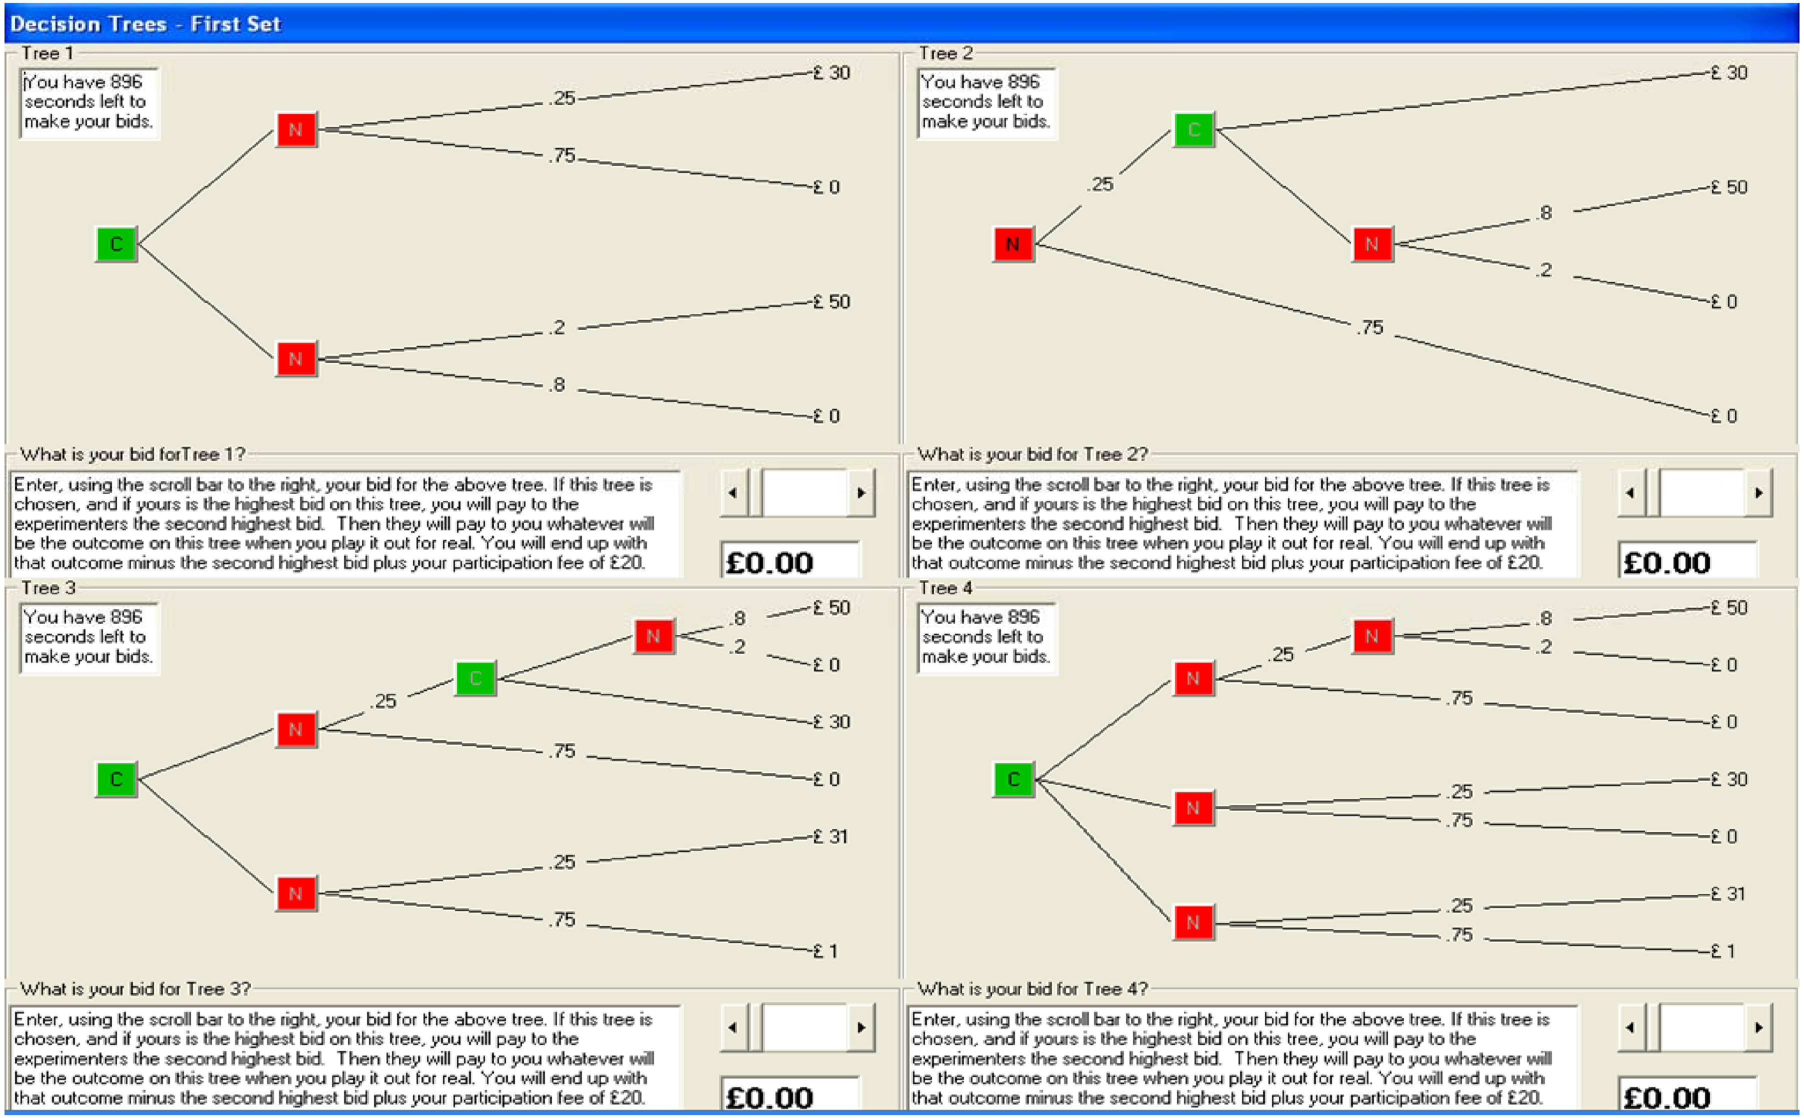

Figure 3: The *ex post* probabilities of the various specifications for all subjects

Red - g parameter significantly different from 1 at 1%
Blue - g parameter significantly different from 1 at 5%
Grey - g parameter insignificantly different from 1
Black - g parameter insignificantly different from 1 but GAUSS unable to calculate standard errors.

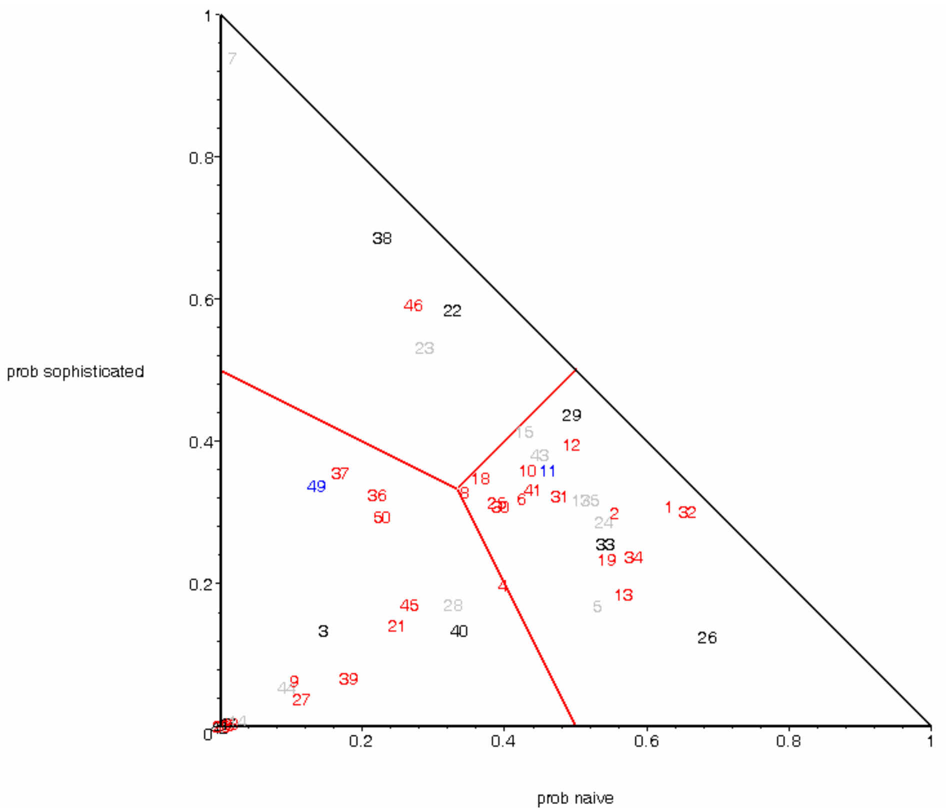

Figure 4: The *ex post* probabilities for those subjects who appear to have non-EU preferences

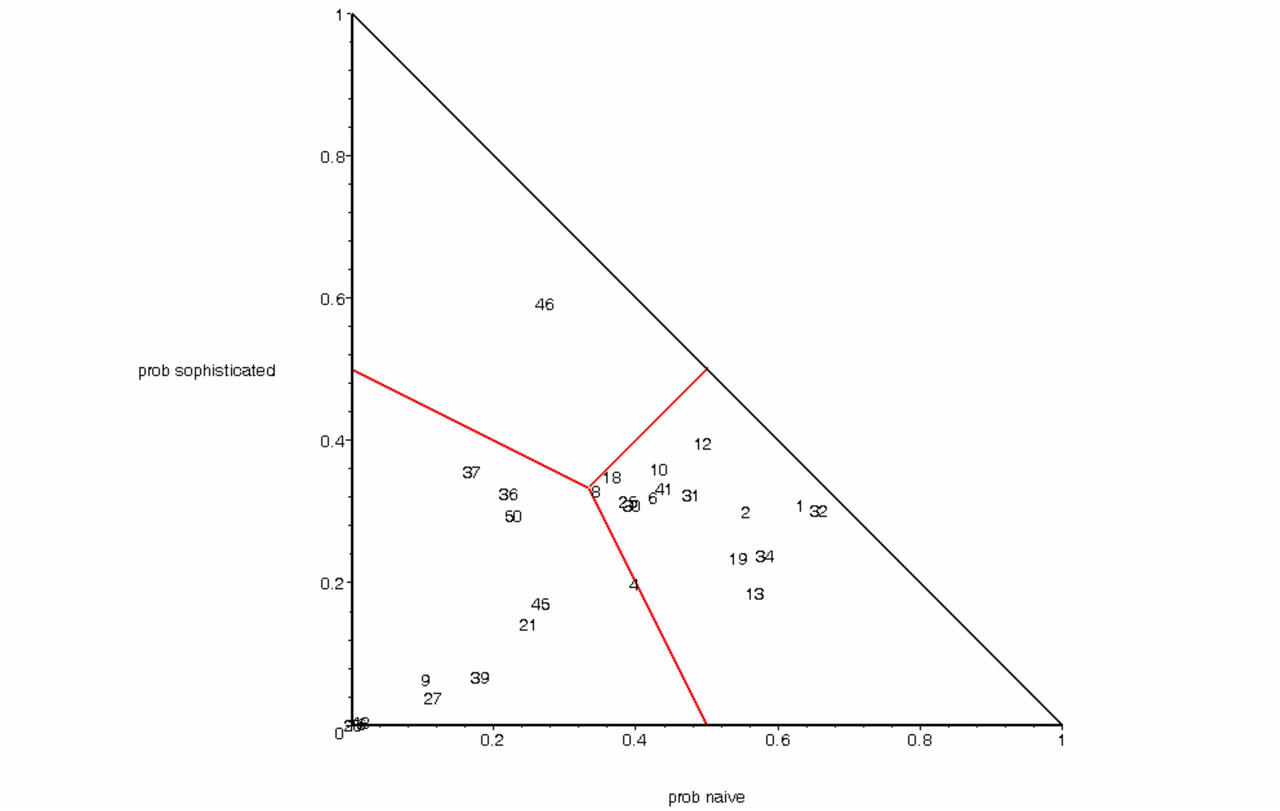

Technical Appendix 1: Detail concerning the construction of the likelihood functions

This appendix derives all the material for the maximum likelihood estimation. We assume that all monetary evaluations of gambles and certainties are done with an error *e* that has an extreme value distribution with mean 0 and parameter *s,* so that the density function of the error is

$$f(e) = \exp(-\exp(-es))\exp(-es)s \tag{7}$$

It follows that the *cdf* is given by

$$F(e) = \exp(-\exp(-es)) \tag{8}$$

Hence the monetary evaluation *X* of a gamble or a certainty with true certainty equivalent *m* has an extreme value distribution with parameters *m* and *s* and hence with *pdf* given by

$$f(x) = \exp(-\exp(-(x-m)s))\exp(-(x-m)s)s \tag{9}$$

and hence the *cdf* of a gamble with true certainty equivalent *m* is given by

$$F(x) = \exp(-\exp(-(x-m)s)) \tag{10}$$

We note that the mean of a variable with an extreme value distribution with parameters *m* and *s* is *m* + γ/s where γ is the Euler-Mascheroni constant = 0.5772156649.

Now suppose that X_1 and X_2 are two gamble evaluations with true certainty equivalents m_1 and m_2 respectively. We need to work out the probability that X_1 is greater than X_2. This is given by

$$\begin{aligned} P(X_1 > X_2) &= \int_{-\infty}^{\infty} F_2(x_1) f_1(x_1) dx_1 \\ &= \int_{-\infty}^{\infty} \exp(-\exp(-(x_1-m_2)s))\exp(-\exp(-(x_1-m_1)s))\exp(-(x_1-m_1)s)s dx_1 \end{aligned} \tag{11}$$

This we can write as

$$\begin{aligned} &P(X_1 > X_2) = \\ &= \int_{-\infty}^{\infty} \exp[-\exp(-(x_1-m_2)s)-\exp(-(x_1-m_1)s)]\exp(-(x_1-m_1)s)s dx_1 \end{aligned} \tag{12}$$

This, after some algebraic manipulation, simplifies to

$$P(X_1 > X_2) = \frac{e^{m_1 s}}{(e^{m_1 s} + e^{m_2 s})} \tag{13}$$

Let us denote this by $x_{2(1)}(m_1, m_2, s)$ and let $x_{2(2)}(m_1, m_2, s)$ denote the residual probability.

We hence have that

$$
\begin{aligned}
x_{2(1)}(m_1,m_2,s) &= P(X_1 > X_2) = \frac{e^{m_1 s}}{(e^{m_1 s}+e^{m_2 s})} \\
x_{2(2)}(m_1,m_2,s) &= P(X_2 > X_1) = \frac{e^{m_2 s}}{(e^{m_1 s}+e^{m_2 s})}
\end{aligned}
\tag{14}
$$

Following the same logic for three valuations X_1, X_2 and X_3 (all with extreme value distributions with respective parameters m_1, m_2 and m_3 and s), and using the appropriate notation we get (see also Beggs *et al* (1981)):

$$
\begin{aligned}
x_{3(1)}(m_1,m_2,m_3,s) &= P(X_1 \text{ is max}) = \frac{e^{m_1 s}}{(e^{m_1 s}+e^{m_2 s}+e^{m_3 s})} \\
x_{3(2)}(m_1,m_2,m_3,s) &= P(X_2 \text{ is max}) = \frac{e^{m_2 s}}{(e^{m_1 s}+e^{m_2 s}+e^{m_3 s})} \\
x_{3(3)}(m_1,m_2,m_3,s) &= P(X_3 \text{ is max}) = \frac{e^{m_3 s}}{(e^{m_1 s}+e^{m_2 s}+e^{m_3 s})}
\end{aligned}
\tag{15}
$$

We will also need the probability density function of the maximum of two, and of the maximum of three variables, all with extreme value distributions. Start with the cumulative distribution function of the maximum of two extreme value distributions, and denote this by $H_2(x,m_1,m_2)$. We have

$$
\begin{aligned}
H_2(x,m_1,m_2) &= \Pr(X_1 < x \text{ and } X_2 < x) \\
&= F_1(x)F_2(x) \\
&= \exp(-\exp(-(x-m_1)s))\exp(-\exp(-(x-m_2)s)) \\
&= \exp\{-[\exp(-(x-m_1)s)+\exp(-(x-m_2)s)]\}
\end{aligned}
\tag{16}
$$

It follows that the *pdf* of the maximum of two variables with extreme value distributions is given by

$$
h_2(x,m_1,m_2) = \exp\{-[\exp(-(x-m_1)s)+\exp(-(x-m_2)s)]\}[\exp(-(x-m_1)s)+\exp(-(x-m_2)s)]s \tag{17}
$$

Similarly the *pdf* of the maximum of three variables with extreme value distributions is given by

$$
\begin{aligned}
h_3(x,m_1,m_2,m_3) &= \exp\{-[\exp(-(x-m_1)s)+\exp(-(x-m_2)s)+\exp(-(x-m_3)s)]\} \text{ multiplied by} \\
&\quad [\exp(-(x-m_1)s)+\exp(-(x-m_2)s)+\exp(-(x-m_3)s)]s
\end{aligned}
\tag{18}
$$

We are now in a position to start writing down the various terms in the likelihood function. We use the following notation where the outcomes *a* through *e* vary from set to set (as in Table 1):

K: the certainty of *c*

L: gamble *(a, e; q, 1-q)*

M: gamble *(c, e; r, 1-r)*

N: gamble *(b, d; r, 1-r)*

O: gamble *(a, e; rq, 1-rq)*

P: gamble *((a,e;q,1-q),e;r,1-r)* – this is a compound non-reduced gamble (used for the analysis of the Type 2 sophisticated agents).

Using RDEU the monetary values of these lotteries are

$$VK = c \tag{19}$$

$$VL = u^{-1}(u(e) + w(q)(u(a) - u(e))) \tag{20}$$

$$VM = u^{-1}(u(e) + w(r)(u(c) - u(e))) \tag{21}$$

$$VN = u^{-1}(u(d) + w(r)(u(b) - u(d))) \tag{22}$$

$$VO = u^{-1}(u(e) + w(rq)(u(a) - u(e)) \tag{23}$$

$$VP = u^{-1}(u(e) + w(r)w(q)(u(a) - u(e)) \tag{24}$$

We note that:

In **Tree 1** there is just one choice: U or D at the only decision node. We code these 1 and 2 respectively and denote the decision by *F*.

In **Tree 2** there is just one decision: U or D at the only decision node. We code these 1 and 2 respectively and denote the decision by *G*.

In **Tree 3** there are two decision nodes. At the first we code the decisions 1 (Up) and 2 (Down) and denote the decision by *H*. At the second we code the decisions 1 (Up) and 2 (Down) and denote the decision by *I*. In the estimation these are combined into one variable *HI* given as follows:

HI = 1 if *H = 1* and *I = 1*

HI = 2 if *H = 1* and *I = 2*

HI = 3 if *H = 2*.

In **Tree 4** there is just one decision: Up, Middle or Down. We denote the decision by *J* and code the respective responses *1, 2* and *3*.

1 TREE 1

1.1 Decisions

This is straightforward and all types of decision makers do the same thing. We have

F = 1 implies the gamble *(c, e; r, 1-r)* - this is Gamble *M*.

F = 2 implies the gamble *(a, e; rq, 1-rq)* - this is Gamble *O*.

We assume that all the valuations are made with error – following the extreme value distribution.

Using our results above we have the following for all types of subject:

$$\begin{aligned} P(F=1) &= x_{2(1)}(VM, VO, s) = \frac{e^{VMs}}{e^{VMs}+e^{VOs}} \\ P(F=2) &= x_{2(2)}(VM, VO, s) = \frac{e^{VOs}}{e^{VMs}+e^{VOs}} \end{aligned} \tag{25}$$

1.2 Bids

These are the same for all types of subject. The bid is given by $B1 = \max(VM, VO)$

The distribution of the right-hand side has *pdf* given by $h_1(VM,VO)$ and hence the probability density at *B1* is given by

$$\begin{aligned} &\exp\{-[\exp(-(B1-VM)s)+\exp(-(B1-VO)s)]\} \text{ times} \\ &[\exp(-(B1-VM)s)+\exp(-(B1-VO)s)]s \end{aligned} \tag{26}$$

This is the probability density associated with the bid B1 for Tree 1 for all subjects, whether they are naïve, resolute, Type 1 or Type 2 sophisticated.

2 TREE 2

2.1 Naive

2.1.1 Decisions

The naïve subject takes the decision at the decision node and argues as follows:

G = 1 implies the certainty of *c* – this is gamble *K*.

G = 2 implies the gamble *(a, e; q, 1-q)* - this is Gamble *L*.

Hence we have:

$$\begin{aligned} P(G=1) &= x_{2(1)}(VK, VL, s) = \frac{e^{VKs}}{e^{VKs}+e^{VLs}} \\ P(G=2) &= x_{2(2)}(VK, VL, s) = \frac{e^{VLs}}{e^{VKs}+e^{VLs}} \end{aligned} \tag{27}$$

2.1.2 Bids

The naïve subject bases his or her bids as viewed from the beginning of the tree. As viewed from the beginning the two choices appear to be the same as in Tree 1. We therefore have that the probability density of the bid is given by

$$\begin{aligned} &\exp\{-[\exp(-(B2-VM)s)+\exp(-(B2-VO)s)]\} \text{ times} \\ &[\exp(-(B2-VM)s)+\exp(-(B2-VO)s)]s \end{aligned} \tag{28}$$

This is the probability density associated with the bid B2 for Tree 2 for a naive person.

2.2 Resolute

2.2.1 Decisions

A resolute person takes the decision at the beginning of the tree, where the choice is exactly the same as in Tree 1. Thus for a resolute person we have:

$$\begin{aligned} P(G=1) &= x_{2(1)}(VM, VO, s) = \frac{e^{VMs}}{e^{VMs}+e^{VOs}} \\ P(G=2) &= x_{2(2)}(VM, VO, s) = \frac{e^{VOs}}{e^{VMs}+e^{VOs}} \end{aligned} \tag{29}$$

2.2.2 Bids

Tree 2 looks exactly like Tree 1 to a resolute person. So the *pdf* of the bid is given by

$$\begin{aligned}&\exp\{-[\exp(-(B2-VM)s)+\exp(-(B2-VO)s)]\}\text{ times}\\&[\exp(-(B2-VM)s)+\exp(-(B2-VO)s)]s\end{aligned}\tag{30}$$

2.3 Type 1 Sophisticated

2.3.1 Decisions

The decision is the same as for the Naive. The Type 1 sophisticated subject takes the decision at the decision node and argues as follows:

G = 1 implies the certainty of *c* – this is gamble *K*.

G = 2 implies the gamble *(a, e; q, 1-q)* - this is Gamble *L*.

Hence we have:

$$\begin{aligned}P(G=1)&=x_{2(1)}(VK,VL,s)=\frac{e^{VKs}}{e^{VKs}+e^{VLs}}\\P(G=2)&=x_{2(2)}(VK,VL,s)=\frac{e^{VLs}}{e^{VKs}+e^{VLs}}\end{aligned}\tag{31}$$

2.3.2 Bids

A Type 1 sophisticated subject works backward using reduction. Then if *K* is valued more than *L*, then, after reduction, Tree 2 gives the lottery *M*, while if *K* is valued less than *L* then, after reduction, Tree 2 gives the lottery *O*. The sophisticated subject values *K* and *L* at the beginning of the tree and decides on the bid on the basis of that decision. If he or she plans to play *K* then the tree is equivalent to *M*, and if he or she plans to play *L* then the tree is equivalent to *O*. So the bid is based on *VM* with probability *P(G=1)* and is based on *VO* with probability *P(G=2)*.

Thus the *pdf* of the bid is given by a weighted average of that based on *VM* and that based on *VO*:

$$\begin{aligned}&\frac{e^{VKs}}{e^{VKs}+e^{VLs}}\exp(-\exp(-(B2-VM)s))\exp(-(B2-VM)s)s\\&+\frac{e^{VLs}}{e^{VKs}+e^{VLs}}\exp(-\exp(-(B2-VO)s))\exp(-(B2-VO)s)s\end{aligned}\tag{32}$$

This is the *pdf* of the bid for Tree 2 for a Type 1 sophisticated subject.

2.4 Type 2 Sophisticated

We note that the only difference between Type 1 and Type 2 is in reduction. Type 2 have lottery *P* in trees 2, 3 and 4 instead of Lottery *O*. To a RDEU person lotteries *O* and *P* are not the same.

2.4.1 Decisions

The decision is the same as for the Naive. The Type 2 sophisticated subject takes the decision at the decision node and argues as follows:

G = 1 implies the certainty of *c* – this is gamble *K*.

G = 2 implies the gamble *(a, e; q, 1-q)* - this is Gamble *L*.

Hence we have:

$$\begin{aligned} P(G=1) = x_{2(1)}(VK, VL, s) = \frac{e^{VKs}}{e^{VKs}+e^{VLs}} \\ P(G=2) = x_{2(2)}(VK, VL, s) = \frac{e^{VLs}}{e^{VKs}+e^{VLs}} \end{aligned} \tag{33}$$

2.4.2 Bids

A Type 2 sophisticated subject backwardly inducts using certainty equivalents – he or she evaluates the trees by substituting in certainty equivalents. Then if *K* is valued more than *L*, after substitution, Tree 2 gives the lottery *M*, while if *K* is valued less than *L* then, after substitution, Tree 2 gives the lottery *P*. The sophisticated subject values *K* and *L* at the beginning of the tree and decides on the bid on the basis of that decision. If he or she plans to play *K* then the tree is equivalent to *M*, and if he or she plans to play *L* then the tree is equivalent to *P.* So the bid is based on *VM* with probability *P(G=1)* and is based on *VP* with probability *P(G=2)*.

Thus the *pdf* of the bid is given by a weighted average of that based on *VM* and that based on *VP*:

$$\begin{aligned} &\frac{e^{VKs}}{e^{VKs}+e^{VLs}} \exp(-\exp(-(B2-VM)s)) \exp(-(B2-VM)s)s \\ &+\frac{e^{VLs}}{e^{VKs}+e^{VLs}} \exp(-\exp(-(B2-VP)s)) \exp(-(B2-VP)s)s \end{aligned} \tag{34}$$

This is the *pdf* of the bid for Tree 2 for a Type 2 sophisticated subject.

3 TREE 3

3.1 Naive

3.1.1 Decisions

The decision at the second node is determined as follows:

I = 1 implies the gamble *(a, e; q, 1-q)* - this is Gamble *L*.

I = 2 implies the certainty of *c* – this is gamble *K*.

Hence we have:

$$
\begin{aligned}
P(I=1) &= x_{2(1)}(VL, VK, s) = \frac{e^{VLs}}{e^{VKs}+e^{VLs}} \\
P(I=2) &= x_{2(2)}(VL, VK, s) = \frac{e^{VKs}}{e^{VKs}+e^{VLs}}
\end{aligned}
\tag{35}
$$

Note here that Up (*I = 1)* is not the same as Up in Tree 2 (*G = 1)* because *L* and *K* are inverted.

The decision at the first node is taken on the basis of how things look at that point. The naïve agent simply looks at the options open at the first node, *O, M* and *N*, and chooses on the basis of them. So we get:

$$
\begin{aligned}
P(H=1) &= x_{3(1)}(VO, VM, VN, s) + x_{3(2)}(VO, VM, VN) = \frac{e^{VOs}+e^{VMs}}{e^{VOs}+e^{VMs}+e^{VNs}} \\
P(H=2) &= x_{3(3)}(VO, VM, VN, s) = \frac{e^{VNs}}{e^{VOs}+e^{VMs}+e^{VNs}}
\end{aligned}
\tag{36}
$$

Putting these together we get

$$
\begin{aligned}
P(HI=1) &= P(H=1)P(I=1) = \frac{e^{VOs}+e^{VMs}}{e^{VOs}+e^{VMs}+e^{VNs}} \frac{e^{VLs}}{e^{VKs}+e^{VLs}} \\
P(HI=2) &= P(H=1)P(I=2) = \frac{e^{VOs}+e^{VMs}}{e^{VOs}+e^{VMs}+e^{VNs}} \frac{e^{VKs}}{e^{VKs}+e^{VLs}} \\
P(HI=3) &= P(H=2) = \frac{e^{VNs}}{e^{VOs}+e^{VMs}+e^{VNs}}
\end{aligned}
\tag{37}
$$

3.1.2 Bids

The bid of a naïve person is simply given by $B3 = \max(VO, VM, VN)$

The distribution of the right-hand side has *pdf* given by $h_3(VO,VM,VN)$ and hence the probability density of the bid is given by

$$\begin{aligned}&\exp\{-[\exp(-(B3-VO)s)+\exp(-(B3-VM)s)+\exp(-(B3-VN)s)]\} \text{ times}\\&[\exp(-(B3-VO)s)+\exp(-(B3-VM)s)+\exp(-(B3-VN)s)]s\end{aligned} \tag{38}$$

This is the probability density associated with the bid B3 for Tree 3 for a naive person.

3.2 Resolute

3.2.1 Decisions

A resolute person evaluates the various strategies from the beginning.

HI = 1 implies *H=1* and *I=1* which implies the gamble *(a, e; rq, 1-rq)* - this is Gamble *O*.

HI = 2 implies *H=1* and *I=2* which implies the gamble *(c, e; r, 1-r)* – this is Gamble *M*.

HI = 3 implies *H=2* which implies the gamble *(b, d; r, 1-r)* - this is Gamble *N*.

We thus have

$$\begin{aligned}P(HI=1)&=x_{3(1)}(VO,VM,VN,s)=\frac{e^{VOs}}{e^{VOs}+e^{VMs}+e^{VNs}}\\P(HI=2)&=x_{3(2)}(VO,VM,VN,s)=\frac{e^{VMs}}{e^{VOs}+e^{VMs}+e^{VNs}}\\P(HI=3)&=x_{3(3)}(VO,VM,VN,s)=\frac{e^{VNs}}{e^{VOs}+e^{VMs}+e^{VNs}}\end{aligned} \tag{39}$$

3.2.2 Bids

His or her bid is given by $B3=\max(VO,VM,VN)$

The distribution of the right-hand side has *pdf* given by $h_3(VO,VM,VN)$ and hence the probability density of the bid is given by

$$\begin{aligned}&\exp\{-[\exp(-(B3-VO)s)+\exp(-(B3-VM)s)+\exp(-(B3-VN)s)]\} \text{ times}\\&[\exp(-(B3-VO)s)+\exp(-(B3-VM)s)+\exp(-(B3-VN)s)]s\end{aligned} \tag{40}$$

This is the probability density associated with the bid B3 for Tree 3 for a resolute person.

3.3 Type 1 Sophisticated

3.3.1 Decisions

Consider first the second decision node.

I = 1 implies the gamble *(a, e; q, 1-q)* - this is Gamble *L*.

$I = 2$ implies the certainty of c - this is Gamble K.

Using our results above we have the following:

$$\begin{aligned} P(I=1) &= x_{2(1)}(VL, VK, s) = \frac{e^{VLs}}{e^{VKs} + e^{VLs}} \\ P(I=2) &= x_{2(2)}(VL, VK, s) = \frac{e^{VKs}}{e^{VKs} + e^{VLs}} \end{aligned} \tag{41}$$

Now go back to the first decision node. A Type 1 sophisticated person will anticipate their future decision. They decide this in advance. If the decision is to play L then moving Up at the first decision node gives them the lottery O; if the decision is to play K then moving Up at the first node gives them the lottery M.

So if L is valued higher than K the individual's decision at the first node is based on VO and VN; whereas if K is valued higher than L the individual's decision at the first node is based on VM and VN. So if $VL > VK$, that is $I = 1$, we have

$$\begin{aligned} P(H=1 \mid I=1) &= x_{2(1)}(VO, VN, s) = \frac{e^{VOs}}{e^{VOs} + e^{VNs}} \\ P(H=2 \mid I=1) &= x_{2(2)}(VO, VN, s) = \frac{e^{VNs}}{e^{VOs} + e^{VNs}} \end{aligned} \tag{42}$$

whereas if $VK > VL$, that is $I = 2$, we have

$$\begin{aligned} P(H=1 \mid I=2) &= x_{2(1)}(VM, VN, s) = \frac{e^{VMs}}{e^{VMs} + e^{VNs}} \\ P(H=2 \mid I=2) &= x_{2(2)}(VM, VN, s) = \frac{e^{VNs}}{e^{VMs} + e^{VNs}} \end{aligned} \tag{43}$$

Putting all the results in equations (44), (45) and (46) together, we have that

$$\begin{aligned} P(HI=1) &= P(H=1 \mid I=1)P(I=1) = \frac{e^{VOs}}{e^{VOs} + e^{VNs}} \frac{e^{VLs}}{e^{VKs} + e^{VLs}} \\ P(HI=2) &= P(H=1 \mid I=2)P(I=2) = \frac{e^{VMs}}{e^{VMs} + e^{VNs}} \frac{e^{VKs}}{e^{VKs} + e^{VLs}} \\ P(HI=3) &= P(H=2 \mid I=1)P(I=1) + P(H=2 \mid I=2)P(I=2) \\ &= \frac{e^{VNs}}{e^{VOs} + e^{VNs}} \frac{e^{VLs}}{e^{VKs} + e^{VLs}} + \frac{e^{VNs}}{e^{VMs} + e^{VNs}} \frac{e^{VKs}}{e^{VKs} + e^{VLs}} \end{aligned} \tag{47}$$

3.3.2 Bids

A Type 1 sophisticated subject works backward using reduction. Then if *K* is valued more than *L*, then, after reduction, Tree 3 gives a choice between the lottery *M* and the lottery *N*; while if *K* is valued less than *L* then, after reduction, Tree 3 gives a choice between the lottery *O* and the lottery *N*. The sophisticated subject values *K* and *L* at the beginning of the tree and decides on the bid on the basis of that decision. If he or she plans to play *K* then the tree is equivalent to the best of *M* and *N*, and if he or she plans to play *L* then the tree is equivalent to the best of *O* and *N.* So the bid is based on the best of *VM* and *VN* with probability *P(I=2)* and is based on the best of *VO* and *VN* with probability *P(I=1)*.

Thus the *pdf* of the bid is given by the following weighted average:

$$\begin{aligned}&\frac{e^{VKs}}{e^{VKs}+e^{VLs}}\exp\{-[\exp(-(B3-VM)s)+\exp(-(B3-VN)s)]\}[\exp(-(B3-VM)s)+\exp(-(B3-VN)s)]s\\&+\frac{e^{VLs}}{e^{VKs}+e^{VLs}}\exp\{-[\exp(-(B3-VO)s)+\exp(-(B3-VN)s)]\}[\exp(-(B3-VO)s)+\exp(-(B3-VN)s)]s\end{aligned} \qquad (48)$$

This is the *pdf* of the bid for Tree 3 for a Type 1 sophisticated subject.

3.4 Type 2 Sophisticated

In essence, these are exactly the same as for the Type 1 sophisticated except that the lottery O is replaced everywhere by the lottery P. However, for completeness, we provide the detail here.

3.4.1 Decisions

Consider first the second decision node.

I = 1 implies the gamble *(a, e; q, 1-q)* - this is Gamble *L*.

I = 2 implies the certainty of *c* - this is Gamble *K*.

Using our results above we have the following:

$$\begin{aligned}P(I=1)&=x_{2(1)}(VL,VK,s)=\frac{e^{VLs}}{e^{VKs}+e^{VLs}}\\P(I=2)&=x_{2(2)}(VL,VK,s)=\frac{e^{VKs}}{e^{VKs}+e^{VLs}}\end{aligned} \qquad (49)$$

Now go back to the first decision node. A Type 2 sophisticated person will anticipate their future decision. They decide this in advance. If the decision is to play *L* then moving Up at the first decision node gives them the lottery *P*; if the decision is to play *K* then moving Up at the first node gives them the lottery *M*.

So if *L* is valued higher than *K* the individual's decision at the first node is based on *VP* and *VN*; whereas if *K* is valued higher than *L* the individual's decision at the first node is based on *VM* and *VN*. So if *VL > VK*, that is *I = 1*, we have

$$
\begin{aligned}
P(H=1\,|\,I=1) &= x_{2(1)}(VP,VN,s) = \frac{e^{VPs}}{e^{VPs}+e^{VNs}} \\
P(H=2\,|\,I=1) &= x_{2(2)}(VP,VN,s) = \frac{e^{VNs}}{e^{VPs}+e^{VNs}}
\end{aligned}
\tag{50}
$$

whereas if *VK > VL*, that is *I = 2*, we have

$$
\begin{aligned}
P(H=1\,|\,I=2) &= x_{2(1)}(VM,VN,s) = \frac{e^{VMs}}{e^{VMs}+e^{VNs}} \\
P(H=2\,|\,I=2) &= x_{2(2)}(VM,VN,s) = \frac{e^{VNs}}{e^{VMs}+e^{VNs}}
\end{aligned}
\tag{51}
$$

Putting all the results in equations (52), (53) and (54) together, we have that

$$
\begin{aligned}
&P(HI=1) = P(H=1\,|\,I=1)P(I=1) = \frac{e^{VPs}}{e^{VPs}+e^{VNs}}\frac{e^{VLs}}{e^{VKs}+e^{VLs}} \\
&P(HI=2) = P(H=1\,|\,I=2)P(I=2) = \frac{e^{VMs}}{e^{VMs}+e^{VNs}}\frac{e^{VKs}}{e^{VKs}+e^{VLs}} \\
&P(HI=3) = P(H=2\,|\,I=1)P(I=1) + P(H=2\,|\,I=2)P(I=2) \\
&= \frac{e^{VNs}}{e^{VPs}+e^{VNs}}\frac{e^{VLs}}{e^{VKs}+e^{VLs}} + \frac{e^{VNs}}{e^{VMs}+e^{VNs}}\frac{e^{VKs}}{e^{VKs}+e^{VLs}}
\end{aligned}
\tag{55}
$$

3.4.2 Bids

A Type 2 sophisticated subject works backwards using certainty equivalents. Then if *K* is valued more than *L*, then, after reduction, Tree 3 gives a choice between the lottery *M* and the lottery *N*; while if *K* is valued less than *L* then, after reduction, Tree 3 gives a choice between the lottery *P* and the lottery *N*. The sophisticated subject values *K* and *L* at the beginning of the tree and decides on the bid on the basis of that decision. If he or she plans to play *K* then the tree is equivalent to the best of *M* and *N*, and if he or she plans to play *L* then the tree is equivalent to the best of *P* and *N*. So the bid is based on the best of *VM* and *VN* with probability *P(I=2)* and is based on the best of *VP* and *VN* with probability *P(I=1)*.

Thus the *pdf* of the bid is given by the following weighted average:

$$\begin{aligned}&\frac{e^{VKs}}{e^{VKs}+e^{VLs}}\exp\{-[\exp(-(B3-VM)s)+\exp(-(B3-VN)s)]\}[\exp(-(B3-VM)s)+\exp(-(B3-VN)s)]s\\&+\frac{e^{VLs}}{e^{VKs}+e^{VLs}}\exp\{-[\exp(-(B3-VP)s)+\exp(-(B3-VN)s)]\}[\exp(-(B3-VP)s)+\exp(-(B3-VN)s)]s\end{aligned} \tag{56}$$

This is the *pdf* of the bid for Tree 3 for a Type 2 sophisticated subject.

4 TREE 4

4.1 Decisions

In Tree 4 naïve, resolute and Type 1 sophisticated subjects follow the same rule. We have

J = 1 implies the gamble *(a, e; rq, 1-rq)* - this is Gamble *O*.

J= 2 implies the gamble *(c, e; r, 1-r)* - this is Gamble *M*.

J = 3 implies the gamble *(b, d; r, 1-r)* - this is Gamble *N*.

Using our results above we have the following for these types of subjects:

$$\begin{aligned}P(J=1)&=x_{3(1)}(VO,VM,VN,s)=\frac{e^{VOs}}{e^{VOs}+e^{VMs}+e^{VNs}}\\P(J=2)&=x_{3(2)}(VO,VM,VN,s)=\frac{e^{VMs}}{e^{VOs}+e^{VMs}+e^{VNs}}\\P(J=3)&=x_{3(3)}(VO,VM,VN,s)=\frac{e^{VNs}}{e^{VOs}+e^{VMs}+e^{VNs}}\end{aligned} \tag{57}$$

4.2 Bids

For the naïve, resolute and Type 1 sophisticated subjects the bid for Tree 4 is given by $B4=\max(VO,VM,VN)$

The distribution of the right-hand side has *pdf* given by $h_4(VO, VM, VN)$ and hence the probability density of the bid B4 is given by

$$\begin{aligned}&\exp\{-[\exp(-(B4-VO)s)+\exp(-(B4-VM)s)+\exp(-(B4-VN)s)]\} \text{ times}\\&[\exp(-(B4-VO)s)+\exp(-(B4-VM)s)+\exp(-(B4-VN)s)]s\end{aligned} \tag{58}$$

This is the probability density associated with the bid B4 for Tree 4 for all naïve, resolute and Type 1 sophisticated subjects.

4.3 Type 2 Sophisticated

The decisions and the bids are exactly the same as for the other types except that the lottery O is replaced everywhere by the lottery P.

Technical Appendix 2: The GAUSS estimation program for the CARA and Quiggin combination

```
library maxlik;

/* This program estimates an RDEU functional fitted to the Hey and Lotito data. Here I am normalising the utility function to have u(0) = 0 and u(150) = 1, and am estimating the parameter s of the extreme value distribution as specified in Technical Appendix 1. We assume that the errors are on the monetary evaluation of the lotteries and the errors have an Extreme Value Distribution.This particular program uses CARA and the Quiggin weighting function. . */

specf=1; /*spec=1 is rank dependent, spec=2 is expected utility specification. */
specl=1;
firsts=1;  /* this is the first subject you want to analyse. */
lasts=50;   /* this is the last subject you want to analyse.  */
mods=5;    /* this represents the models you want to estimate:
            1 is naive
            2 is resolute
            3 is type 1 sophisticated
            4 is type 2 sophisticated
            5 is all of them */
ssm=0.3;  /* this is the minimum STARTING value for the s parameter. */
ssi=0.1;  /* this is the increment for the STARTING value for the s parameter. */
ssl=0.7;  /* this is the maximum STARTING value for the s parameter. The minimum is ssm and it increases in steps of ssi. */
minr=-0.1; /* this is the minimum value of r */
maxr=0.1; /* this is the maximum absolute value for the r parameter allowed. */
ming=0.25; /* this is the minimum value of the g parameter allowed. */
maxg=25.0; /* this is the maximum value of the g parameter allowed. */
maxs=10.0; /* this is the maximum value of the s parameter allowed. */
mins=0.0; /* this is the minimun value of the s parameter allowed.  */
/* gam=0.5772156649;  this is the gamma parameter of the theory.  */
gam=0;  /* I am planning edit out all reference to the gamma */
df=0.01; /* this is to convert the bid data from pence to pounds. */
n=50;    /* this is the number of subjects. */
k=24;    /* this is the number of observations for each subject. */
_max_MaxIters = 1000; /*this is the maximum number of iterations allowed. */

output file = "d:/active/experiments/hey and lotito/first experiment/GAUSS estimation programs/formulation 3/rdaq1.out";
format /rd 10,5;

o3=ones(3,1);
z3=zeros(3,1);
dt1=ones(3,1);
dt2=ones(3,1);
dt3=ones(3,1);
dt4=ones(3,1);
bt1=ones(3,1);
bt2=ones(3,1);
bt3=ones(3,1);
bt4=ones(3,1);

a=50|150|60;
b=31|51|41;
c=30|50|40;
d=1|1|1;
e=0|0|0;
p=1.0|1.0|1.0;
q=0.8|0.5|0.75;
r=0.25|0.1|0.2;
pars=ones(3,1);
```

```
spec=specf;
do while spec<=specl;
output on;
if spec==1;
_max_active=1|1|1;
print "rank dependent specification";
endif;
if spec==2;
_max_active=1|0|1;
print "expected utility specification";
endif;
output off;
vars=1;
vvv=ones(n,1);
load data[n,k+3] = "d:/active/experiments/hey and lotito/first experiment/GAUSS estimation programs/formulation 3/alldata.out";
output on;screen off;
print "cAra and the Quiggin weighting function";
print "m    subject    s start    retcode    ll    r    g    s    raw r    raw g    raw s  sd(raw r)  sd(raw g)  sd(raw s)  ";
output off;screen on;

sn=firsts;

do while sn<=lasts;
bt1[1]=df*data[sn,1];bt1[2]=df*data[sn,5];bt1[3]=df*data[sn,9];
bt2[1]=df*data[sn,2];bt2[2]=df*data[sn,6];bt2[3]=df*data[sn,10];
bt3[1]=df*data[sn,3];bt3[2]=df*data[sn,7];bt3[3]=df*data[sn,11];
bt4[1]=df*data[sn,4];bt4[2]=df*data[sn,8];bt4[3]=df*data[sn,12];
dt1[1]=data[sn,13];dt1[2]=data[sn,17];dt1[3]=data[sn,21];
dt2[1]=data[sn,14];dt2[2]=data[sn,18];dt2[3]=data[sn,22];
dt3[1]=data[sn,15];dt3[2]=data[sn,19];dt3[3]=data[sn,23];
dt4[1]=data[sn,16];dt4[2]=data[sn,20];dt4[3]=data[sn,24];

if mods==1 or mods==5;
ss=ssm;
do while ss<=ssl;
start= ircf(0.0)|igcf(1.0)|iscf(ss);
{xn,ffn,gn,cn,rn} = maxlik(vvv,vars,&lln,start);
if rn>=0 and ffn*24>-100;
output on;
print "n " sn~ss~rn~ffn*24~rcf(xn[1])~gcf(xn[2])~scf(xn[3])~xn'~sqrt(diag(cn))';
output off;
endif;
ss=ss+ssi;
endo;
endif;

if mods==2 or mods==5;
ss=ssm;
do while ss<=ssl;
start= ircf(0.0)|igcf(1.0)|iscf(ss);
{xr,fr,gr,cr,rr} = maxlik(vvv,vars,&llr,start);
if rr>=0 and fr*24>-100;
output on;
print "r " sn~ss~rr~fr*24~rcf(xr[1])~gcf(xr[2])~scf(xr[3])~xr'~sqrt(diag(cr))';
output off;
endif;
ss=ss+ssi;
endo;
endif;
```

```
if mods==3 or mods==5;
ss=ssm;
do while ss<=ssl;
start= ircf(0.0)|igcf(1.0)|iscf(ss);
{xs,fs,gs,cs,rs} = maxlik(vvv,vars,&lls,start);
if rs>=0 and fs*24>-100;
output on;
print "s1" sn~ss~rs~fs*24~rcf(xs[1])~gcf(xs[2])~scf(xs[3])~xs'~sqrt(diag(cs))';
output off;
endif;
ss=ss+ssi;
endo;
endif;

if mods==4 or mods==5;
ss=ssm;
do while ss<=ssl;
start= ircf(0.0)|igcf(1.0)|iscf(ss);
{xc,fc,gc,cc,rc} = maxlik(vvv,vars,&llc,start);
if rc>=0 and fc*24>-100;
output on;
print "s2" sn~ss~rc~fc*24~rcf(xc[1])~gcf(xc[2])~scf(xc[3])~xc'~sqrt(diag(cc))';
output off;
endif;
ss=ss+ssi;
endo;
endif;

sn=sn+1;
endo;
spec=spec+1;
endo;

proc lln(x,y);
/* log-likelihood for the naive subjects*/
local ll,llb1,llb2,llb3,llb4,lld1,lld2,lld3,lld4;
local llb1r,llb2r,llb3r,llb4r;
local d11,d12;
local d21,d22;
local d31,d32,d33;
local d41,d42,d43;
local rp,gp,sp;
local VK,VL,VM,VN,VO;
local ua,ub,uc,ud,ue;

rp=rcf(x[1]);
gp=gcf(x[2]);
sp=scf(x[3]);

ua=u(a,rp);
ub=u(b,rp);
uc=u(c,rp);
ud=u(d,rp);
ue=u(e,rp);

VK=c;
VL=ui(ue+w(q,gp).*(ua-ue),rp);
VM=ui(ue+w(r.*p,gp).*(uc-ue),rp);
VN=ui(ud+w(r.*p,gp).*(ub-ud),rp);
VO=ui(ue+w(r.*q,gp).*(ua-ue),rp);

d11=VM*sp-ln(exp(VM*sp)+exp(VO*sp));
d12=VO*sp-ln(exp(VM*sp)+exp(VO*sp));
```

```
lld1=(2*o3-dt1).*d11+(dt1-o3).*d12;

d21=VK*sp-ln(exp(VK*sp)+exp(VL*sp));
d22=VL*sp-ln(exp(VK*sp)+exp(VL*sp));
lld2=(2*o3-dt2).*d21+(dt2-o3).*d22;

d31=VL*sp+ln(exp(VO*sp)+exp(VM*sp))-ln(exp(VK*sp)+exp(VL*sp))-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d32=VK*sp+ln(exp(VO*sp)+exp(VM*sp))-ln(exp(VK*sp)+exp(VL*sp))-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d33=VN*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
lld3=0.5*(3*o3-dt3).*(2*o3-dt3).*d31;
lld3=lld3+(3*o3-dt3).*(dt3-o3).*d32;
lld3=lld3+0.5*(dt3-2*o3).*(dt3-o3).*d33;

d41=VO*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d42=VM*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d43=VN*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
lld4=0.5*(3*o3-dt4).*(2*o3-dt4).*d41+(3*o3-dt4).*(dt4-o3).*d42+0.5*(dt4-2*o3).*(dt4-o3).*d43;

llb1r=exp(-(bt1-VM)*sp-gam*o3)+exp(-(bt1-VO)*sp-gam*o3);
llb1=-llb1r+ln(llb1r)+o3*ln(sp);

llb2r=exp(-(bt2-VM)*sp-gam*o3)+exp(-(bt2-VO)*sp-gam*o3);
llb2=-llb2r+ln(llb2r)+o3*ln(sp);

llb3r=exp(-(bt3-VO)*sp-gam*o3)+exp(-(bt3-VM)*sp-gam*o3)+exp(-(bt3-VN)*sp-gam*o3);
llb3=-llb3r+ln(llb3r)+o3*ln(sp);

llb4r=exp(-(bt4-VO)*sp-gam*o3)+exp(-(bt4-VM)*sp-gam*o3)+exp(-(bt4-VN)*sp-gam*o3);
llb4=-llb4r+ln(llb4r)+o3*ln(sp);

ll=o3'*(llb1+llb2+llb3+llb4+lld1+lld2+lld3+lld4);

retp(ll);
endp;


proc llr(x,y);
/* the log-likelihood for the resolute subjects */
local ll,llb1,llb2,llb3,llb4,lld1,lld2,lld3,lld4;
local llb1r,llb2r,llb3r,llb4r;
local d11,d12;
local d21,d22;
local d31,d32,d33;
local d41,d42,d43;
local rp,gp,sp;
local VK,VL,VM,VN,VO;
local ua,ub,uc,ud,ue;

rp=rcf(x[1]);
gp=gcf(x[2]);
sp=scf(x[3]);

ua=u(a,rp);
ub=u(b,rp);
uc=u(c,rp);
ud=u(d,rp);
ue=u(e,rp);

VK=c;
VL=ui(ue+w(q,gp).*(ua-ue),rp);
VM=ui(ue+w(r.*p,gp).*(uc-ue),rp);
VN=ui(ud+w(r.*p,gp).*(ub-ud),rp);
VO=ui(ue+w(r.*q,gp).*(ua-ue),rp);
```

```
d11=VM*sp;
d12=VO*sp;
lld1=(2*o3-dt1).*d11+(dt1-o3).*d12-ln(exp(VM*sp)+exp(VO*sp));

d21=VM*sp-ln(exp(VM*sp)+exp(VO*sp));
d22=VO*sp-ln(exp(VM*sp)+exp(VO*sp));
lld2=(2*o3-dt2).*d21+(dt2-o3).*d22;

d31=VO*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d32=VM*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d33=VN*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
lld3=0.5*(3*o3-dt3).*(2*o3-dt3).*d31+(3*o3-dt3).*(dt3-o3).*d32+0.5*(dt3-2*o3).*(dt3-o3).*d33;


d41=VO*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d42=VM*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d43=VN*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
lld4=0.5*(3*o3-dt4).*(2*o3-dt4).*d41+(3*o3-dt4).*(dt4-o3).*d42+0.5*(dt4-2*o3).*(dt4-o3).*d43;

llb1r=exp(-(bt1-VM)*sp-gam*o3)+exp(-(bt1-VO)*sp-gam*o3);
llb1=-llb1r+ln(llb1r)+o3*ln(sp);

llb2r=exp(-(bt2-VM)*sp-gam*o3)+exp(-(bt2-VO)*sp-gam*o3);
llb2=-llb2r+ln(llb2r)+o3*ln(sp);

llb3r=exp(-(bt3-VO)*sp-gam*o3)+exp(-(bt3-VM)*sp-gam*o3)+exp(-(bt3-VN)*sp-gam*o3);
llb3=-llb3r+ln(llb3r)+o3*ln(sp);

llb4r=exp(-(bt4-VO)*sp-gam*o3)+exp(-(bt4-VM)*sp-gam*o3)+exp(-(bt4-VN)*sp-gam*o3);
llb4=-llb4r+ln(llb4r)+o3*ln(sp);

ll=o3'*(llb1+llb2+llb3+llb4+lld1+lld2+lld3+lld4);

retp(ll);
endp;



proc lls(x,y);
/* the log-likelihood for the type 1 sophisticated subjects*/
local ll,llb1,llb2,llb3,llb4,lld1,lld2,lld3,lld4;
local llb1r,llb2r,llb3r,llb4r;
local llb21,llb22,llb31,llb32;
local d11,d12;
local d21,d22;
local d31,d32,d33,d331,d332,d333,d334;
local d41,d42,d43;
local rp,gp,sp;
local VK,VL,VM,VN,VO;
local ua,ub,uc,ud,ue;

rp=rcf(x[1]);
gp=gcf(x[2]);
sp=scf(x[3]);

ua=u(a,rp);
ub=u(b,rp);
uc=u(c,rp);
ud=u(d,rp);
ue=u(e,rp);

VK=c;
VL=ui(ue+w(q,gp).*(ua-ue),rp);
```

```
VM=ui(ue+w(r.*p,gp).*(uc-ue),rp);
VN=ui(ud+w(r.*p,gp).*(ub-ud),rp);
VO=ui(ue+w(r.*q,gp).*(ua-ue),rp);

d11=VM*sp-ln(exp(VM*sp)+exp(VO*sp));
d12=VO*sp-ln(exp(VM*sp)+exp(VO*sp));
lld1=(2*o3-dt1).*d11+(dt1-o3).*d12;

d21=VK*sp-ln(exp(VK*sp)+exp(VL*sp));
d22=VL*sp-ln(exp(VK*sp)+exp(VL*sp));
lld2=(2*o3-dt2).*d21+(dt2-o3).*d22;

d31=VO*sp+VL*sp-ln(exp(VO*sp)+exp(VN*sp))-ln(exp(VK*sp)+exp(VL*sp));
d32=VM*sp+VK*sp-ln(exp(VM*sp)+exp(VN*sp))-ln(exp(VK*sp)+exp(VL*sp));
d331=exp(VN*sp)./(exp(VO*sp)+exp(VN*sp));
d332=exp(VL*sp)./(exp(VK*sp)+exp(VL*sp));
d333=exp(VN*sp)./(exp(VM*sp)+exp(VN*sp));
d334=exp(VK*sp)./(exp(VK*sp)+exp(VL*sp));
d33=ln(d331.*d332+d333.*d334);
lld3=0.5*(3*o3-dt3).*(2*o3-dt3).*d31+(3*o3-dt3).*(dt3-o3).*d32+0.5*(dt3-2*o3).*(dt3-o3).*d33;

d41=VO*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d42=VM*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
d43=VN*sp-ln(exp(VO*sp)+exp(VM*sp)+exp(VN*sp));
lld4=0.5*(3*o3-dt4).*(2*o3-dt4).*d41+(3*o3-dt4).*(dt4-o3).*d42+0.5*(dt4-2*o3).*(dt4-o3).*d43;

llb1r=exp(-(bt1-VM)*sp-gam*o3)+exp(-(bt1-VO)*sp-gam*o3);
llb1=-llb1r+ln(llb1r)+o3*ln(sp);

llb21=exp(-(bt2-VM)*sp-gam*o3);
llb21=exp(-llb21).*llb21;
llb21=llb21.*(exp(VK*sp)./(exp(VK*sp)+exp(VL*sp)));
llb22=exp(-(bt2-VO)*sp-gam*o3);
llb22=exp(-llb22).*llb22;
llb22=llb22.*(exp(VL*sp)./(exp(VK*sp)+exp(VL*sp)));
llb2=ln(llb21+llb22)+o3*ln(sp);

llb31=exp(-(bt3-VM)*sp-gam*o3)+exp(-(bt3-VN)*sp-gam*o3);
llb31=exp(-llb31).*llb31;
llb31=llb31.*(exp(VK*sp)./(exp(VK*sp)+exp(VL*sp)));
llb32=exp(-(bt3-VO)*sp-gam*o3)+exp(-(bt3-VN)*sp-gam*o3);
llb32=exp(-llb32).*llb32;
llb32=llb32.*(exp(VL*sp)./(exp(VK*sp)+exp(VL*sp)));
llb3=ln(llb31+llb32)+o3*ln(sp);

llb4r=exp(-(bt4-VO)*sp-gam*o3)+exp(-(bt4-VM)*sp-gam*o3)+exp(-(bt4-VN)*sp-gam*o3);
llb4=-llb4r+ln(llb4r)+o3*ln(sp);

ll=o3'*(llb1+llb2+llb3+llb4+lld1+lld2+lld3+lld4);

retp(ll);
endp;

proc llc(x,y);
/* the log-likelihood for the type 2 sophisticated subjects*/
local ll,llb1,llb2,llb3,llb4,lld1,lld2,lld3,lld4;
local llb1r,llb2r,llb3r,llb4r;
local llb21,llb22,llb31,llb32;
local d11,d12;
local d21,d22;
local d31,d32,d33,d331,d332,d333,d334;
local d41,d42,d43;
local rp,gp,sp;
```

```
local VK,VL,VM,VN,VO,VP;
local ua,ub,uc,ud,ue;

rp=rcf(x[1]);
gp=gcf(x[2]);
sp=scf(x[3]);

ua=u(a,rp);
ub=u(b,rp);
uc=u(c,rp);
ud=u(d,rp);
ue=u(e,rp);

VK=c;
VL=ui(ue+w(q,gp).*(ua-ue),rp);
VM=ui(ue+w(r.*p,gp).*(uc-ue),rp);
VN=ui(ud+w(r.*p,gp).*(ub-ud),rp);
VO=ui(ue+w(r.*q,gp).*(ua-ue),rp);
VP=ui(ue+w(r,gp).*w(q,gp).*(ua-ue),rp);

d11=VM*sp-ln(exp(VM*sp)+exp(VO*sp));
d12=VO*sp-ln(exp(VM*sp)+exp(VO*sp));
lld1=(2*o3-dt1).*d11+(dt1-o3).*d12;

d21=VK*sp-ln(exp(VK*sp)+exp(VL*sp));
d22=VL*sp-ln(exp(VK*sp)+exp(VL*sp));
lld2=(2*o3-dt2).*d21+(dt2-o3).*d22;

d31=VP*sp+VL*sp-ln(exp(VP*sp)+exp(VN*sp))-ln(exp(VK*sp)+exp(VL*sp));
d32=VM*sp+VK*sp-ln(exp(VM*sp)+exp(VN*sp))-ln(exp(VK*sp)+exp(VL*sp));
d331=exp(VN*sp)./(exp(VP*sp)+exp(VN*sp));
d332=exp(VL*sp)./(exp(VK*sp)+exp(VL*sp));
d333=exp(VN*sp)./(exp(VM*sp)+exp(VN*sp));
d334=exp(VK*sp)./(exp(VK*sp)+exp(VL*sp));
d33=ln(d331.*d332+d333.*d334);
lld3=0.5*(3*o3-dt3).*(2*o3-dt3).*d31+(3*o3-dt3).*(dt3-o3).*d32+0.5*(dt3-2*o3).*(dt3-o3).*d33;

d41=VO*sp-ln(exp(VP*sp)+exp(VM*sp)+exp(VN*sp));
d42=VM*sp-ln(exp(VP*sp)+exp(VM*sp)+exp(VN*sp));
d43=VN*sp-ln(exp(VP*sp)+exp(VM*sp)+exp(VN*sp));
lld4=0.5*(3*o3-dt4).*(2*o3-dt4).*d41+(3*o3-dt4).*(dt4-o3).*d42+0.5*(dt4-2*o3).*(dt4-o3).*d43;

llb1r=exp(-(bt1-VM)*sp-gam*o3)+exp(-(bt1-VO)*sp-gam*o3);
llb1=-llb1r+ln(llb1r)+o3*ln(sp);

llb21=exp(-(bt2-VM)*sp-gam*o3);
llb21=exp(-llb21).*llb21;
llb21=llb21.*(exp(VK*sp)./(exp(VK*sp)+exp(VL*sp)));
llb22=exp(-(bt2-VP)*sp-gam*o3);
llb22=exp(-llb22).*llb22;
llb22=llb22.*(exp(VL*sp)./(exp(VK*sp)+exp(VL*sp)));
llb2=ln(llb21+llb22)+o3*ln(sp);

llb31=exp(-(bt3-VM)*sp-gam*o3)+exp(-(bt3-VN)*sp-gam*o3);
llb31=exp(-llb31).*llb31;
llb31=llb31.*(exp(VK*sp)./(exp(VK*sp)+exp(VL*sp)));
llb32=exp(-(bt3-VP)*sp-gam*o3)+exp(-(bt3-VN)*sp-gam*o3);
llb32=exp(-llb32).*llb32;
llb32=llb32.*(exp(VL*sp)./(exp(VK*sp)+exp(VL*sp)));
llb3=ln(llb31+llb32)+o3*ln(sp);

llb4r=exp(-(bt4-VP)*sp-gam*o3)+exp(-(bt4-VM)*sp-gam*o3)+exp(-(bt4-VN)*sp-gam*o3);
llb4=-llb4r+ln(llb4r)+o3*ln(sp);
```

```
ll=o3'*(llb1+llb2+llb3+llb4+lld1+lld2+lld3+lld4);

retp(ll);
endp;

proc u(x,r);
local uv;
if r==0;
uv=x/150;
endif;
if not r==0;
uv=(o3-exp(-r*x))/(1-exp(-150*r));
endif;
retp(uv);
endp;

proc ud(x,r);
local udv;
if r==0;
udv=1/150;
endif;
if not r==0;
udv=r*exp(-r*x)/(1-exp(-150*r));
endif;
retp(udv);
endp;

proc ui(u,r);
local xv,uv,nr,nc,i,j,onc;
nr=rows(u);
nc=cols(u);
onc=ones(nr,nc);
i=1;
do while i<=nr;
j=1;
do while j<=nc;
if u[i,j]>1;u[i,j]=1;endif;
j=j+1;
endo;
i=i+1;
endo;
if r==0;
xv=150*u;
endif;
if not r==0;
xv=-ln(onc-u*(1-exp(-150*r)))/r;
endif;
retp(xv);
endp;

proc w(p,g);
local wv,i,j,nr,nc;
nr=rows(p);
nc=cols(p);
wv=zeros(nr,nc);
i=1;
do while i<=nr;
j=1;
do while j<=nc;
if p[i,j]<=0;wv[i,j]=0;endif;
if p[i,j]>0 and p[i,j]<1;wv[i,j]=(p[i,j]^g)/((p[i,j]^g)+(1-p[i,j])^g)^(1/g);endif;
if p[i,j]>=1;wv[i,j]=1;endif;
```

```
j=j+1;
endo;
i=i+1;
endo;
retp(wv);
endp;

proc rcf(x);
local cv;
x=real(x);
cv=(maxr-minr)/(1+exp(-x))+minr;
retp(cv);
endp;

proc ircf(r);
local xv;
xv=-ln((maxr-r)/(r-minr));
retp(xv);
endp;

proc gcf(x);
local cv;
x=real(x);
cv=(maxg-ming)/(1+exp(-x))+ming;
retp(cv);
endp;

proc igcf(g);
local xv;
xv=-ln((maxg-g)/(g-ming));
retp(xv);
endp;

proc scf(x);
local cv;
x=real(x);
cv=(maxs-mins)/(1+exp(-x))+mins;
retp(cv);
endp;

proc iscf(s);
local sv;
sv=-ln((maxs-s)/(s-mins));
retp(sv);
endp;

closeall;
end.
```

Technical Appendix 3: Estimates for subjects 1, 2 and 3 for the CARA Quiggin combination

Subject 1	loglik	prob	r par	g par	s par	raw r	raw g	raw s	sd(raw r)	sd(raw g)	sd(raw s)
n	-12.71074	0.3652	0.03298	0.42831**	0.93140	0.68516	-4.92584	-2.27589	0.11278	0.27567	0.22672
r	-14.14725	0.0868	0.03636	0.45751**	0.86856	0.76195	-4.77300	-2.35265	0.13545	0.30174	0.22734
s1	-13.35608	0.1915	0.03449	0.46092**	0.84282	0.71919	-4.75655	-2.38554	0.12213	0.29491	0.22751
s2	-12.73517	0.3564	0.03063	0.43496**	0.83134	0.63281	-4.88891	-2.40051	0.12896	0.31714	0.22246

Subject 2	loglik	prob	r par	g par	s par	raw r	raw g	raw s	sd(raw r)	sd(raw g)	sd(raw s)
n	-16.96547	0.4551	-0.04365	2.89288**	0.51578	-0.93590	-2.12403	-2.91171	0.34123	0.22652	0.25110
r	-18.31379	0.1182	-0.04411	2.87034**	0.54075	-0.94720	-2.13362	-2.86180	0.29764	0.20270	0.24932
s1	-17.89108	0.1804	-0.04200	2.78074**	0.52210	-0.89534	-2.17245	-2.89885	0.27758	0.19498	0.25284
s2	-17.57957	0.2463	-0.03843	2.73582**	0.55937	-0.81026	-2.1923	-2.82596	0.22483	0.16499	0.25478

Subject 3	loglik	prob	r par	g par	s par	raw r	raw g	raw s	sd(raw r)	sd(raw g)	sd(raw s)
n	-15.62822	0.1291	-0.00963	3.22367**	0.73770	-0.19329	-1.99103	-2.53017	0.12232	0.21796	0.22631
r	-14.03244	0.6365	0.01512	9.40525!!	0.85589	0.30479	-0.53261	-2.36873	9999.9999	9999.9999	9999.9999
s1	-15.69680	0.1205	-0.00659	2.93079**	0.72087	-0.13205	-2.10807	-2.55507	0.12080	0.20124	0.24279
s2	-15.75262	0.1140	-0.00744	3.01370**	0.71165	-0.14915	-2.07385	-2.56892	0.11848	0.19960	0.24000

Glossary of columns

loglik: maximised log-likelihood

prob: Bayesian posterior probability of specification (taking all four specifications into account)

r par: estimate of *r*

g par: estimate of *g:* ** (*) indicates significantly different from 1 at 1% (5%) level; !! indicates that maxlik could not calculate standard error of *g*.

s par: estimate of *s*

raw r: is the estimate of the transformed *r* parameter

raw g: is the estimate of the transformed *g* parameter

raw s: is the estimate of the transformed *r* parameter

sd(raw r): is the standard error of the estimate of the transformed *r* parameter

sd(raw g): is the standard error of the estimate of the transformed *g* parameter
sd(raw e): is the standard error of the estimate of the transformed *s* parameter
"9999.9999" indicates that Maxlik was unable to calculate the standard error.
Note: to resolve problems of convergence of the GAUSS maxlik routine we transformed the parameters so that the estimates would not go outside reasonable ranges. The remaining columns are the estimates and standard errors of the estimates of the transformed parameters. In each case the transformation was of the form: $y = f(x) = (max-min)/(1+exp(-x))+min$ and the inverse $x = f^{-1}(y) = - ln((max-y)/(y-min))$. The *min* and *max* varied from parameter to parameter. For the CARA Quiggin case we have: for *r*: $min = -0.1$ and $max=0.1$; for *g*: $min = 0.25$ and $max = 25.0$; and for *s*: $min = 0.0$ and max =10.0.

Technical Appendix 4: Instructions

EXEC

INSTRUCTIONS

Preamble

Welcome to this experiment. It is an experiment on the economics of dynamic decision making under risk. The Ministry for Education, University and Research of Italy (MIUR) has provided the funds to finance this research. Thank you for taking part.

Please read these instructions carefully. It is important that you do so, as your payment for taking part in this experiment will depend wholly or in part on the decisions that you take. The payment will be made, in cash, at the conclusion of the experiment. The payment will consist of a participation fee of £20, plus or minus whatever extra money you make or lose as a result of the decisions you make during the experiment. Unless you choose otherwise, you can not walk away with any less than the participation fee of £20. However, depending upon your decisions, those of the other participants, and chance, you could walk away with considerably more or considerably less. You will be asked to sign a receipt for the payment, and to acknowledge that you participated voluntarily in the experiment. The results of the experiment will be used for the purpose of academic research and will be published in such a way that your anonymity will be preserved.

The Experiment

In the experiment you will be presented with three sets of four *decision problems* – that is a total of 12 decision problems. The purpose of the experiment is to discover your *willingness to pay* for each of these decision problems. They are described in detail at the end of these instructions. Each decision problem has at most five possible outcomes which differ for each set: £50, £31, £30, £1 and £0 for the first set, £150, £51, £50, £1 and £0 for the second set and £60, £41, £40, £1 and £0 for the third set, the specific outcome depending upon chance. In each decision problem you will be told exactly how the outcome is determined. You will be in a group with four other participants and we will be selling off a randomly chosen *one* of these problems to one of the five participants in your group. You are invited to *bid* for each of the four problems in the three sets.

Bidding

The way that you bid is simple. You will all be sat at individual computer terminals and the computer will explain to you what you should do. You will be allowed to practice playing out all the decision problems while you are deciding on your bids. You will type in your bids for each of the four problems in the three sets. You will have plenty of time to practice playing out with the problems and decide on your bids – to be precise 20 minutes for each set of decision problems. After you have typed all your bids, the computer will show all your bids for all the three sets of decision problems.

Practicing playing out the decision problems

During the bidding period you will be allowed to practice playing out the decision problems as much as you want. The way that you practice playing out the problems is exactly the same way as you will play them out for real. This is described in detail below. Of course, the outcomes when you practice playing out the decision problems do not affect your payment in any way.

Playing out the decision problems for real

After all participants have entered all their bids, you will play out all the decision problems for real, taking whatever decision you prefer. Your outcomes for each of the decision problems will then be displayed on your screen.

How we provide an incentive for you to reveal your willingness to pay through your bids

We do this by the way that we will sell off one of the decision problems to one of the participants. After all 5 members of the group have entered their bids, we will do the following. We will display on the computer screens the bids of the 5 members of the group for all the three sets of decision problems. The five participants will chose one of their 5 members to pick **one** of the 12 problems at random. We will identify the member of the group with the highest bid for that problem and that highest bidder will pay to us the bid of the second highest bidder. We will pay to the highest bidder the outcome which resulted from their playing out for real. Note also that if the highest bidder pays more than £20 (the participation fee) plus whatever outcome resulted, he or she will walk away from the experiment with less money than he or she came with. We give below three examples.

Examples

Example 1: the highest bidder bids £8, the second highest bidder bids £7 and the outcome to the highest bidder is £40. In this case, the highest bidder gets £20 (the participation fee) *minus* £7 (the second highest bid) *plus* £40 (the outcome) – that is, £20 - £7 + £40 = £53. All the other participants get £20 (the participation fee).

Example 2: the highest bidder bids £32, the second highest bidder bids £30 and the outcome to the highest bidder is £60. In this case, the highest bidder gets £20 (the participation fee) *minus* £30 (the second highest bid) *plus* £60 (the outcome) – that is, £20 - £30 + £60 = £50. All the other participants get £20 (the participation fee).

Example 3: the highest bidder bids £35, the second highest bidder bids £33 and the outcome to the highest bidder is £0. In this case, the highest bidder gets £20 (the participation fee) *minus* £33 (the second highest bid) *plus* £0 (the outcome) – that is, £20 - £33 + £0 = -£13. That is, the highest bidder *loses* £13. All the other participants get £20 (the participation fee).

Note crucially

To avoid problems caused by a participant losing money and not having the resources to pay, we will require the highest bidder to pay to us the second highest bid. Therefore if you do not have the cash with you to fund any possible losses, you should modify your bids appropriately. Note that if you bid too high, you could walk away from this experiment with less money than when you came. Of course, you can avoid this by the simple strategy of never bidding more than the participation fee of £20, though you should also note that this might mean that you do not walk away from the experiment with more than £20.

What you should do

It should be clear that your bid for each of the problems should be equal to your willingness to pay for that problem. If you bid more than that, you may end up paying more for the problem than you think it is worth; if you bid less, you may end up not getting the problem even though you would have been willing to pay more for it. As this point is important, we repeat it here: **your bid for each of the decision problems should be equal to your willingness to pay for that problem.** The purpose of the experiment is to discover your willingness to pay for each of the decision problems.

The decision problems

The decision problems will be presented to you in the form shown in the attached figure. Here we are just showing you the first set of four decision problems. The decision problems in the other two sets are identical in structure, and differ only in the payoffs and probabilities. You will work through the problems from the left box, finally ending up with one of the payoffs on the right. The green boxes indicate Choice nodes, where you take a decision. The red boxes indicate moves by Nature – Nature moves probabilistically with the probabilities indicated on the lines to the right of the red box.

Problem 1 (Tree 1)
You first click on the green box (C for Choice); two boxes will appear: U for Up and D for Down: if you click on U, you will move to the top red box (N for Nature) and when you click on it, Nature will move in such a way that you have either 25% probability of receiving £30 or 75% probability of receiving £0; if you click on D, you will move to the bottom red box (N for Nature) and when you click on it, you have either 20% probability of receiving £50 or 80% probability of receiving £0.

Problem 2 (Tree 2)
You start by clicking on the first red box (N for Nature). You have either 75% probability of receiving £0 or 25% probability of getting to the green box (C for Choice). If you get there you should click on this green box and two boxes will appear: U for Up and D for Down: if you click on U, you will get £30; if you click on D, you will end up at the second red box (N for Nature) and when you click on it, you have either 80% probability of receiving £50 or 20% probability of receiving £0.

Problem 3 (Tree 3)
You first click on the first green box (C for Choice); two boxes will appear: U for Up and D for Down: if you click on U, you will move to the top red box (N for Nature) and when you click on it, you have either 75% probability of receiving £0 or 25% probability of ending up at the second green box (C for Choice), and if so, you should then click on the second green box and two boxes

will appear: U for Up and D for Down; if you click D at the second green box, you will get £30; if you click U at the second green box, you will end up at the red box (N for Nature): when you click on it, you have either 80% probability of receiving £50 or 20% probability of receiving £0. In case you choose to click on D at the first green box, you will end up at the bottom red box (N for Nature): when you click on it, you have either 25% probability of receiving £31 or 75% probability of receiving £1.

Problem 4 (Tree 4)
You will have to click on the green box (C for Choice); three boxes will appear: 1, 2 and 3.
If you click 1, you will end up at the top red box (N for Nature): when you click on it, you have either 75% probability of receiving £0 or 25% probability of ending up at the second top red box (N for Nature). When you click on the second red box (N for Nature), you have either 80% probability of receiving £50 or 20% probability of receiving £0. If you click 2, you will end up at the middle red box (N for Nature): when you click on it, you have either 25% probability of receiving £30 or 75% probability of receiving £0. If you click 3, you will end up at the bottom red box (N for Nature): when you click on it, you have either 25% probability of receiving £31 or 75% probability of receiving £1.

Note: in the Instructions at this point was a screen shot. This is exactly the same as Figure 2 in the text.