

An experiment on risky choice amongst
households.

by

Ian Bateman

School of Environmental Sciences and CSERGE,

University of East Anglia,

Norwich NR4 7TJ,

UK

and Adjunct, ARE, University of Western Australia

and

Alistair Munro.

(Corresponding author. Email to a.a.munro@uea.ac.uk):

School of Economic and Social Studies,

University of East Anglia, Norwich, NR4 7TJ.

Abstract.

Many important decisions are taken within multi-adult households rather than by individuals. This paper reports the results of the first economic experiment designed to test theories of *household* rather than *individual* choice. We use a sample of established couples and face them individually and jointly with decision tasks involving lotteries. We test whether their choices conform to expected utility theory. We find that choices made by couples exhibit the same kinds of patterns (e.g. the common ratio and common consequence effects) as are regularly recorded with individuals, but that choices made jointly are more risk averse than those made separately.

Keywords: Household choice, experiment, expected utility.

JEL Codes: C920, D130, D80.

Introduction.

The expected utility-maximizing household is one of the most common models employed to understand economic behaviour. This *standard model*, which is used to investigate saving, insurance decisions, labour supply etc., involves two important assumptions. First, that the household acts as if it has a single set of preferences and secondly that these preferences conform to the axioms of expected utility theory (EUT). The first assumption has received a significant amount of scrutiny (e.g. Browning and Chiappori et al, 1998), but very little attention has been paid to the second assumption for households as opposed to individuals. In fact though there is copious experimental evidence on how individuals choose, to date there has been very little experimental investigation into how multi-adult households or couples make their decisions.¹

This paper therefore presents results of an experiment designed to investigate the following issue: to what extent do the decisions made by couples and the decisions made separately by individuals who are part of a couple conform to the standard model? In outline the experiment is as follows: we use a sample of established couples² and present them with tasks involving binary choices between lotteries of the kind depicted in Figure 1. In the first section of the experiment the subjects are separated and face choices separately; in the second section they remain apart and must predict their partner's answers from the first section; in the

¹ For instance in Starmer's (2001) survey of the field of risky choice, there is no discussion of evidence on household as opposed to individual behaviour. There is an interesting body of work by psychologists on this issue (see Corfman and Lehmann, 1983 for example), but the questions asked provide little insight into the applicability of economists' models of choice.

² Meaning that the couple are in a relationship of at least one year's standing and live together.

third section they rejoin their partner and make choices as a couple.³ The tasks are not all repeated in each section, though there is some overlap. Each lottery has possible monetary payoffs for each individual within the couple and these payoffs may be different. A random lottery device is used to provide incentives.

[Figure 1 here.]

It might reasonably be supposed that the results of individual choice experiments should carry over into household decision-making. However, leaving aside the issue of differences in subject pool, the decision-making environment of the household might eliminate certain anomalies commonly observed in individual choice. For instance, with two people scrutinising probabilities rather than one, the kind of editing and framing effects which underlie Kahneman and Tversky's 1979 explanation of the Allais paradox might not apply. Conversely, even if individuals separately have preferences which satisfy EUT, the rule used to aggregate preferences within the household might produce choices for the household at variance with the predictions of EUT. So, the existing theory and experimental evidence on individual choice does not therefore imply much about how households make choices in risky situations and in particular, whether households conform to the standard model.

In the standard model of household choice, the household is assumed to be unitary - that is, the household is modelled as a single agent with a single set of preferences - either

³ Bone et al, 1999, 2000 investigate decision-making in pairs, but in their experiment university students are paired at random whereas we are interested in the behaviour of pre-existing decision-making units – i.e. established couples. Secondly, in their design the pairs of students are given a collective payment and must decide how to divide it. Our lotteries assign payments to particular individuals (and our payment procedures reflect this), though this is not to deny the possibility of bargains being made or anticipated.

because there are no public goods local to the household and all members share the same preferences or because the structure of incentives within the household align individual preferences with those of the decision-maker (as in, for instance, the ‘rotten kid theorem’, Becker 1974). Empirical testing (see Lundberg et al, 1997 or Alderman et al, 1996 for example) gives results largely hostile to the unitary model, particularly its prediction of income pooling (IP) which is the property that household behaviour may respond to changes in aggregate household income but not to who in the household earns that income. This has prompted a large number of alternative household models, but nevertheless IP remains a convenient modelling assumption in many contexts. For this experiment we design a mix of tests of EUT, some of which are conditional on households satisfying IP and some of which are not. One reason for having the conditional tests is that in many empirical situations it may not be possible to observe the sources of income in a household. We wish to see whether any departures from EUT are robust in the sense that they are still observable in the face of variation in the identity of the income recipient.

II Theory.

For simplicity we consider a two-person household. Let agent $i = 1, 2$ receive payment m_{is} in state of the world $s = 1, \dots, S$. A typical lottery p (or q, r or s) is then a vector (p_1, \dots, p_S) . The standard sign, \succeq denotes the weak preference relationship for the household, with strict preference denoted \succ and indifference \sim , constructed in the usual manner.

A household obeys expected utility theory in its joint choices if there exists a strictly increasing function $w(m_{1s}, m_{2s})$ such that the household ranks lotteries according to,

$W(p) \equiv \sum_{s=1}^{s=S} p_s w(m_{1s}, m_{2s})$. In other words $W(p) \geq W(q) \leftrightarrow p \succeq q$. In a similar manner it is

also possible to define utility functions, $W^i(p)$ $i=1,2$, for the two individuals. Note that the relationship between the W^i 's and W depends on the household aggregation rule. So the fact that the household choices conform to EUT does not imply that the W^i 's satisfy the axioms of EUT - or vice versa.

We shall say that the household *income pools* (IP) or that it is an *income pooler* if $w(m_{1s}, m_{2s}) = w(m_{1s}', m_{2s}')$ whenever $m_{1s} + m_{2s} = m_{1s}' + m_{2s}'$ for all s .

Although w has two arguments rather than the one that is typical of individual choice, nevertheless for the household or individual which maximizes $W(\cdot)$, preferences between lotteries should have the familiar properties of EUT. Figure 2 shows a standard unit probability triangle representing lotteries involving three possible values of w : w_1 , w_2 and w_3 , with $w_3 > w_2 > w_1$. In the figure, the solid line connecting a and e is parallel to that between c and d .

[Figure 2 here.]

In the unit probability triangle, EUT predicts that indifference curves are straight, parallel lines. However, individuals frequently fail to conform to the predictions of EUT in a number of ways. Possibly, three of the most robust anomalies (see Starmer, 2001 for a comprehensive survey) are the *common ratio* effect, the *common consequence* effect and failure of the *betweenness* property. In the first, individuals tend to choose the safer option represented by choosing a out of the pair $\{a,e\}$ and then the riskier option d out of the pair $\{c,d\}$. In the second a is chosen out of $\{a,b\}$ and then d is picked from $\{c,d\}$. In the third case, EUT implies that individuals who choose b out of $\{a,b\}$ should choose e out of $\{a,e\}$, whereas often individuals choose b and then a . Indifference curves in the triangle therefore seem to be more like the broken lines depicted in figure 2, than the straight and parallel lines implied by EUT. The experiment is designed to see if couples have similar revealed preferences.

III Experimental Design.

Figure 3 summarises the experimental design. Upon entering the venue, one member of the pair was randomly allocated either a ‘wave’ or ‘triangle’ card, their partner receiving the other card. This allocation of cards was then used to separate pair, one set of partners being taken into a separate room. The first two sections of the experiment were conducted with partners in these separate rooms, pairs then rejoined each other for the final section of the experiment. Throughout the experiment the investigators used a script (available from the authors) and at the start of each section subjects received summary instructions for the tasks to be undertaken in that section.

[Figure 3 about here.]

In section one of the experiment (i.e. where subjects are separated from their partners), subjects faced choices, similar to those depicted in figure one. The description of a typical lottery was composed of three elements: ranges of numbers were shown along the top, underneath which were shown corresponding payoffs for the subject, below which were given corresponding payoffs for their partner.⁴ The numbers along the top corresponded to numbered discs in a bag of one hundred discs shown to the subjects by the experimenters. Subjects were told that, at the end of the experiment, one of the choice tasks would be chosen at random for each couple and played out for real. If this was one of the tasks undertaken in section 1, then subjects would play out the lottery they chose in that task by taking a

⁴ For the joint choice questions, the triangle partner’s payoffs were always shown first. We found no evidence that this order gave triangle partners more or less influence in the joint decisions.

numbered disc from the bag, that number determining the resulting payoffs for them and their partner.⁵

Two questions designed to test understanding were placed at the end of the briefing for the first section. After all subjects had answered these questions satisfactorily, they completed the section 1 questions in their own time⁶. Answer books were then collected.

In section 2, subjects (who remained separated from their partners) were asked to predict their partners' answers in section 1. Before doing so, they were led through the relevant instructions, including those concerning incentives (see below). After completing these prediction tasks, subjects were faced with a short questionnaire which collected demographic data. Once this was completed answer books were collected and subjects rejoined their partners for the final section of the experiment.

In the final section, couples made choices jointly. At the start of this section, couples selected a single, small envelope from a shuffled pile of similar envelopes placed in front of them, but were told not to open it until instructed. No prompts were given as to which partner should make this selection. The full set of envelopes contained lottery ticket numbers for all possible questions from all sections, one lottery ticket being placed inside each envelope. The randomly drawn number inside the envelope selected by a couple determined which question they would play out 'for real' at the end of the experiment.

The subjects were given details of how the payout procedures would operate at the end

⁵ Subjects were not told at that stage about the remaining sections of the experiment.

⁶ An initial version of the design consisted of 10 questions in each section. This was subsequently extended to add a further two questions. Standard statistical tests indicated that there was no significant differences between the data for the questions common to both variants and accordingly these responses were pooled within our analysis.

of the experiment and led through the instructions for section 3. Once all subjects had completed their tasks, we began opening the small envelopes, executing lotteries and making payoffs.

The incentive system was as follows: for lottery ticket numbers from 1 to n , the triangle partner played his or her choice for that section 1 question and the wave partner received £0.50 for each correct prediction (in section 2) of their partners section 1 answers. For numbers between $n+1$ and $2n$, the wave partner played his or her choice from section 1 while the triangle partner was paid £0.50 for each correct prediction (in section 2) of their partners section 1 answers. For numbers from $2n+1$ to $3n$, the couple played their joint choice for that question from section 3 and no money was paid for predictions. This random lottery system is incentive compatible if individuals are selfish⁷ (and make no binding agreements on ex-post trade), but it would be usual to suppose some degree of other-regarding preferences within couples. As a result, it is conceivable that an altruist might view the first two sections of the experiment as an exercise in co-ordination and possibly choose so as to maximize the predictive success of his or her partner. We aimed to guard against this possibility in three ways. First, prior to the experiment we told subjects only that the research was aiming 'to help us understand how couples make decisions'. As a second measure the prediction questions (in section 2) always came after the separate choice questions (in section 1) and we saved the briefing for section 2 until all subjects in a session had completed section 1. So, subjects therefore had no reason to anticipate that they should answer in section 1 so as to raise the possible payoffs of their partner. As a final measure we kept the payments for prediction to a relatively small fraction of the payments associated with the choice sections. The payments

⁷ Cubitt et al 1998 provide evidence that random lottery schemes are a reliable means of eliciting preferences even when subjects are not EUT maximizers.

obtainable from the choice section ranged from £0 to £80, whereas the range for the prediction section was only £0 to £12. And in the choice questions, the expected values of the options always differed by more than £0.50. This means that a risk neutral income pooling agent (for example) would not make an expected gain from switching choices in order to improve the predictive success of their partner.

We went to some trouble to preserve the confidentiality of the answers from sections 1 and 2. Partners were paid sequentially and separately with any payments placed in envelopes. The payment process occurred in another room or in a position which masked any payments made. Subjects were not informed of their partner's answers in section 1 of the experiment and they were not given information about the accuracy of their partner's predictions.⁸ Now, theories of the household are rarely explicit on whether individuals are privy to the patterns of consumption and income of their partners. However, there is plenty of empirical evidence of asymmetric information within the household. For instance, in a survey of spending habits in UK households, Pahl, 1983 reports that 'typically, husbands over-estimated the amounts wives spent on leisure, while wives under-estimated how much their husbands spent', p. 132, (see also Woolley, 2000 and Treas, 1993). It is reasonable to expect theories to be robust in the face of such possibilities. However, our main reason for confidentiality is as follows: Many economic theories of the household relate collective choice to individual preferences over goods. To test such theories we normally require data about individual preferences over commodities. Revealing choices to partners might instead produce information on preferences

⁸ If one partner predicted perfectly (or scored zero), then provided she or he had perfect recall, that subject could deduce a partner's choices. No participant raised this possibility with us during the conduct of the experiment and no-one achieved perfection in their predictions (or scored zero).

over actions. For instance, it might create incentives for individuals to choose so as to garner approval from their partners. Such motives may actually be an important source of household behaviour, but they are not typically the objects of preference in economic theories of the household. So, we opted for confidentiality in our design.

Following a successful pilot session, the experiments were carried out from December 2002 to March 2003. Subjects were recruited from the city of Norwich and rural Norfolk via email, through community groups and using posters. Session sizes varied from two to ten couples and were held at a variety of venues, including a village hall and the experimental economics laboratory at the University of East Anglia. In recruiting we required all individuals to be over 21, to be living with their partners and to have been together as a couple for at least one year. We asked subjects to bring evidence of their relationship and made random checks.⁹

IV Results.

We recruited 76 couples for our experiment. Average payoffs were just under £17 per individual - more than twice the median hourly post-tax wage for a UK adult in 2003. Ages ranged from 22 to 70, with a mean of 37.3. On average couples had been together for 11 years, with a maximum of 46 and a minimum of 1. Seventy-three percent of individuals stated that they were married to their current partner and all the couples in our sample were heterosexual. The distribution of children per couple was bimodal with peaks at zero and two and a mean of 1.1. So, without being representative of the UK adult population, the subjects were generally older and more diverse than the typical sample of university students used in choice

⁹ Evidence included passports, photos, bills to the same address and, in three cases, children.

experiments.

In what follows, the tasks are labelled. Their details can be found in the Appendix. The number identifies the task, while T indicates tasks faced by triangle subjects in their section 1, W stands for tasks faced by wave subjects in their section 1 and J indicates tasks put jointly to couples in the final part of the experiment.¹⁰

For both joint and separate choice we included four questions (T13, J13, W13 and W14) such as the one shown in Figure 1 where one option first-order stochastically dominates the other. For these questions the dominated option was chosen in just under 6% of observations.

Recall that IP is a feature of what we termed the standard model. We had seven tasks where one of the options dominates the other, for subjects whose choices satisfy IP (but not necessarily otherwise). In 90% of cases the choice is in conformity with IP and this accordance is stronger for the choices made jointly than for those made when the individuals are separated. Suppose we hold the null hypothesis that in all cases subjects mean to choose the IP dominating option, but make a mistake in 6% of cases (i.e. the rate of ‘error’ in the choices with one dominating option discussed above). With the exception of one task (T11), the pattern of choices is consistent with this null hypothesis, suggesting that IP is a reasonable assumption in the context of this experiment. We also had a number of tests of IP based on pairs of tasks which are equivalent when faced by a chooser who satisfies the IP property. That data is more mixed in the conclusions it produces.¹¹ So, the evidence for IP in our data is

¹⁰ ‘Groups’ are sets of tasks which are equivalent from the perspective of a chooser who satisfies IP. The task numbers do not match the order of questions, but are purely for reference purposes.

¹¹ For details see Bateman and Munro, 2003, which focuses on the IP issue.

not overwhelming and we therefore conduct tests of EUT both with and without its presence as an auxiliary assumption.

Tables 1-3 summarise tests of EUT - tests in Table 1 are based upon the individual choice data, Table 2 draws upon the prediction tasks and Table 3 is based upon joint choices. All the comparisons shown are within subject or within couple. In the first column of these three tables, entries labelled CR represent common ratio tests, those marked CC represent common consequence tests, while BB indicates tests of the betweenness property of EUT. The next column states whether the comparison is conditional on the assumption of the IP property. If it is, then this means that pairs of tasks can only be plotted in the same unit probability triangle if IP holds. In the two proportions columns, the numbers represent the fraction of the sample choosing the safer option. According to EUT, the fraction should be the same across the relevant tasks. This is always the null hypothesis. According to the typical results of individual choice experiments the proportion in the task 2 column should be lower.¹² This is always the alternative hypothesis. Taking all of the tables together we see only one instance of equality (the EUT prediction) compared to the remaining 32 cases all of which are in the direction of the alternative hypothesis. In the final column we report probability values (to three significant figures) for the null hypothesis that the sample proportions are equal, using a paired, one-sided z-test. For CR, CC and BB, a large number of these comparisons are statistically significant; in many cases at levels of significance well below 0.1%.

[Table 1 here]

¹² Matthew Rabin, 2000, argues that choosing the safe option in choices of this kind is prima facie evidence against the EUT model, because it is incompatible with attitudes to risk displayed in other settings. In our case, the choice between options also typically reflects intrahousehold inequality aversion so the same argument does not automatically apply.

Table 1 summarises results for the cases involving separate choice. All the CR and BB comparisons are statistically significant at the 10% level or lower. For two CC cases where the test is not conditional on IP, the difference in responses to the two tasks is not statistically significant.¹³ So, broadly speaking the evidence for a common ratio effect and for failure of the betweenness property is stronger than that for the common consequence effect. In terms of the stylised indifference curves in Figure 2, it suggests that the section between a and b is roughly parallel to the curve between c and d, but that the indifference curves show increasing risk aversion (i.e. become steeper) between b and e. It is worth noting that this pattern persists across the variety of tasks listed in the table – for some of these tasks all the payoffs are to the choosing agent, but for many tasks both partners might possibly receive payment and in several cases it is only the partner that might receive payments. Nevertheless the pattern of choices is consistent.

Table 2 presents the data from the prediction section of the data. These tasks are the same as those in Table 1, but it is the other partner who is doing the predicting. All the comparisons are statistically significant, even with the CC examples. So, the results suggest that prediction deviates significantly from EUT. When we look at the prediction data in detail we find that partners predict correctly in 65% of cases. This is significantly better than fifty-fifty; it is also better than the success rate if they supposed (as a benchmark example) that their partner was a risk neutral income pooler. However, if individuals predict according to how they themselves choose and preferences are not correlated within couples then the predicted success rate is 64.7% - which is not statistically significantly different from the

¹³ It is also possible to use some Wave responses to create a between-subject CC comparison. When this comparison is made a proportion 0.53 choose the safer option out of {a,b} while 0.43 pick the safer option out of {c,d}. This is significant at the 10% level ($p=0.095$).

actual value.

[Table 2 here].

The fact that, when separated, individual partners depart from the standard model in their choices does not mean that those individuals have non-EUT preferences. They may be altruists who believe that their partners have non-EUT preferences. Similarly, in the absence of common knowledge, it cannot be deduced that individuals who predict anomalous behaviour in their partners actually believe that their partners have non-EUT preferences. If, though, this was the case, then we would expect the possibilities for communication afforded by the joint decision-making responsibility of section 3 to iron out any misunderstandings. In fact, as we can see from Table 3, the joint choice data exhibits the same patterns as the prediction and separate choice data. This is true with and without the auxiliary assumption of IP. Moreover, as with the separated choice, the evidence for common ratio effects and for the failure of the betweenness property is stronger than that for the common consequence effect. This suggests that the departure from EUT observed in the separate choice and predictions data is not simply due to misconceptions about the preferences of partners. Rather, it seems to be a persistent feature of choice in the context of multi-person households.

[Table 3 here].

Table 4 summarises some interesting comparisons of choices made jointly (in Section 3) and when separated (in Section 1) for the four tasks where one option is safer than the other and where all the risk is in one partner's payoffs. Note that, when viewed by Wave and Triangle subjects, these tasks appear *reflected* in the sense that, for any given task, payoffs which belong to the self when Wave chooses belong to the partner when Triangle chooses and vice versa. Two things are particularly notable in Table 4. First, the proportions for Wave and Triangle subjects are very close – in other words subjects appear to place equal weight on

their partner's payoffs as on their own and are not more or less risk averse when it is their partner who faces the risk rather than themselves. Second, perhaps surprisingly, choices made jointly are consistently more risk averse than those made separately, to the extent that in three of the four cases the difference is statistically significant whichever partner is taken as the benchmark. It is not clear to us why this result occurred. Standard risk sharing arguments would predict joint choice would exhibit less risk aversion, since it gives opportunities for agreeing to ex-post risk sharing transfers. We can also rule out misperception of the partner's degree of risk aversion as the explanation, since there is no evidence for such a bias in the prediction data. Possibly, the result is due to the psychology of group choice, one robust feature of which (see Kerr et al, 1996, for instance) is that collective decisions are typically more extreme than their individual counterparts. Yet, in our case it is not clear why the safe option should be viewed as more extreme. A final possibility is suggested by anecdotal evidence from our participants, some of whom suggested a 'fear of recrimination' as a significant factor influencing joint choices. This could make some participants reluctant to be seen to be pressing for the risky option.

[Table 4 here]

V Discussion.

The fact that the preferences of two individuals separately obey the assumptions of EUT does not imply that their collective decisions will always conform to the same axioms. Conversely, depending on the household decision process, it is possible that two individuals with non-EUT preferences can produce collective choices that do satisfy the predictions of EUT. It follows that tests of whether the decisions of established couples conform to EUT are logically separate from the issue of whether individual decisions satisfy the theory. Nevertheless, in this experiment couples show the same anomalous patterns in their risky

choices as have been frequently observed in individual choice experiments. When separated from their partners, individuals who are part of a couple also show the same patterns and predict the same patterns in their partner's choices. The results of the experiment also suggest that the results are robust in the face of changes in the identity of who in the household receives the payoffs.

In the face of individual choice anomalies, a number of alternative theories of risky choice have been put forward (e.g. regret theory, prospect theory, etc.). To a significant degree, these theories have been motivated by ideas drawn from the psychology of the individual. It is not clear that the same ideas automatically apply in the household, where decisions are typically made in an interactive fashion; other forces may be at the root of the results found here. For instance, we found a surprisingly high incidence of examples where, separately both individuals chose the risky option in an identical task, but as couples chose the safe option. If such patterns are a feature of many households it would suggest that behavioural models of collective decision-making may be quite different to their individual decision-making counterparts.

References.

Alderman Harold, Chiappori P.A., Haddad L, Hoddinott J, Kanbur R, 1995, Unitary Versus Collective Models Of The Household - Is It Time To Shift The Burden Of Proof, *World Bank Research Observer* 10 (1): 1-19 Feb.

Bateman, Ian and Alistair Munro, 2003, Testing economic models of the household: an experiment, *CSERGE working paper, University of East Anglia*.

Becker, Gary S., 1974, A theory of social interactions, *Journal of Political Economy*, vol. 82 (6), 1063-1093.

Bergstrom, Theodore, 1989, A fresh look at the rotten kid theorem and other household mysteries, *Journal of Political Economy*, 97(5), 1138-1159.

Bone, John, John Hey and John Suckling, 1999, Are Groups More or Less Consistent Than Individuals?, *Journal of Risk and Uncertainty*, 8, 63-81.

Bone, John, John Hey and John Suckling, 2000, A Simple Risk-Sharing Experiment, *University of York Discussion papers in Economics*, No. 2000/36.

Browning, Martin, and Pierre-Andre Chiappori, 1998. Efficient intra-household allocations: a general characterisation and empirical tests. *Econometrica* 66 (6), 1241–1278.

Corfman, Kim. P. and Donald. R. Lehmann, 1987, Models of cooperative group decision-making and relative influence: An experimental investigation of family purchase decisions. *Journal of Consumer Research*, 14: 1-13.

Cubitt, Robin P., Chris.V. Starmer and Robert Sugden, 1998, On the validity of the random lottery incentive system, *Experimental Economics*, 1, 115-131.

Kahneman, Daniel, and Amos Tversky, 1979, Prospect Theory: an analysis of decision under risk, *Econometrica*, 47: 263-291.

Kerr, Norbert L., MacCoun, Robert J., Kramer, Geoffrey P., 1996 Bias in judgment: Comparing individuals and groups. *Psychological Review*. Oct Vol 103(4) 687-719.

Lundberg, Shelly J., Robert A. Pollak and Terence J. Wales, 1997, Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit, *Journal of Human Resources*, Vol. 32, No. 3. (Summer.), pp. 463-480.

Pahl, Jan, 1990, Household spending, personal spending and the control of money in marriage, *Sociology*, 24:1, 119-138.

Peters, H. Elizabeth, A. Sinan Ünür, Jeremy Clark, William D. Schulze, forthcoming, Free-Riding and the Provision of Public Goods in the Family: An Experimental Test of the Rotten Kid Theorem, *International Economic Review*.

Phipps, Shelley A., and Burton Peter S., 1998, What's Mine is yours? The influence of male and female incomes on patterns of household expenditure, *Economica*, 65, 599-613.

Rabin, Matthew, 2000, Risk Aversion and Expected-Utility Theory: A Calibration Theorem, *Econometrica* 68: 1281-1292.

Treas, Judith, 1993, Money in the Bank: Transaction Costs and the Economic Organization of Marriage, *American Sociological Review* 58: 723-734.

Woolley, Frances, 2000, Control over Money in Marriage, *Carleton Economic Papers* 00-07.

Table 1. Tests of EUT using Separate Choice Data							
N	Type of comparison	IP assumed?	Task 1	Task 2	Proportion choosing safer option,		Probability
					Task 1	Task 2	
76	CR	No	T6	T8	0.66	0.43	0.000***
76	CR	Yes	T6	T7	0.66	0.54	0.073*
76	CR	Yes	T4	T8	0.78	0.43	0.000***
76	CR	Yes	T4	T7	0.78	0.54	0.001***
34	CR	Yes	W4	W7	0.77	0.53	0.001***
34	CR	Yes	W5	W7	0.71	0.53	0.054*
76	CC	No	T1	T8	0.50	0.43	0.190
34	CC	No	W3	W7	0.53	0.53	0.500
34	CC	Yes	T2	T8	0.53	0.32	0.055*
34	CC	Yes	W2	W7	0.71	0.47	0.028**
76	BB	No	T6	T1	0.66	0.50	0.048**
34	BB	Yes	W5	W3	0.71	0.56	0.042**
76	BB	Yes	T4	T1	0.78	0.50	0.000***

*** indicates difference significant at 1% level, 1 tailed test; ** indicates significant at 5% level; * indicates significant at 10% level.

Table 2. EUT and Prediction Data

N	Type of comparison	IP assumed?	Task 1	Task 2	Proportion predicting safer option		Probability
					Task 1	Task 2	
76	CR	No	T6	T8	0.70	0.43	0.000***
76	CR	Yes	T6	T7	0.70	0.59	0.051*
76	CR	Yes	T4	T8	0.88	0.43	0.000***
76	CC	No	T1	T8	0.62	0.43	0.011**
34	CC	No	W3	W7	0.47	0.32	0.078*
34	CC	Yes	T2	T8	0.68	0.41	0.024**
34	CC	Yes	W2	W7	0.68	0.32	0.002***
76	BB	No	T6	T1	0.70	0.62	0.068*
34	BB	Yes	W5	W3	0.71	0.47	0.009***
76	BB	Yes	T4	T1	0.88	0.62	0.000***

*** indicates difference significant at 1% level, 1 tailed test; ** indicates significant at 5% level; * indicates significant at 10% level.

Table 3. Tests of EUT using Joint Choice Data

N	Type of comparison	IP assumed?	Task 1	Task 2	Proportion choosing safer option,		Probability
					Task 1	Task 2	
34	CR	Yes	J4	J7	0.94	0.59	0.000***
34	CR	Yes	J5	J7	0.88	0.59	0.000***
34	CR	No	J10	J7	0.73	0.59	0.013**
34	CC	No	J1	J7	0.71	0.59	0.110
34	CC	Yes	J2	J7	0.71	0.59	0.130
34	CC	Yes	J9	J11	0.44	0.06	0.000***
76	BB	No	J4	J1	0.94	0.64	0.000***
34	BB	Yes	J5	J1	0.88	0.71	0.006***
34	BB	Yes	J4	J2	0.94	0.71	0.002***
34	BB	Yes	J5	J2	0.88	0.71	0.016**

*** indicates difference significant at 1% level, 1 tailed test; ** indicates significant at 5% level; * indicates significant at 10% level.

Table 4. Comparison between choices made jointly and when separated

Task	Proportion choosing safer option			Probability
	Separate choice		Joint Choice	
	Triangle	Wave	Joint	
1	0.5	0.53	0.64	0.030**
4	0.78	0.77	0.94	0.000***
7	0.54	0.53	0.64	0.059*
9	0.35	0.395	0.44	0.364

Probability is the p-value associated with the test of equality between the joint choice value and the closest value from separate choice. *** indicates difference significant at 1% level, 1 tailed test; ** indicates significant at 5% level; * indicates significant at 10% level.

Question 2

Option A

For numbers:	1-50	51-100
You receive	£20	£0
Your partner receives	£0	£20

I choose (*tick one*): Option A

Option B

For numbers:	1-50	51-100
You receive	£20	£40
Your partner receives	£0	£20

Option B

Figure 1. A typical question from section 1 of the experiment.




Figure 2. Tests of EUT in the unit probability triangle.

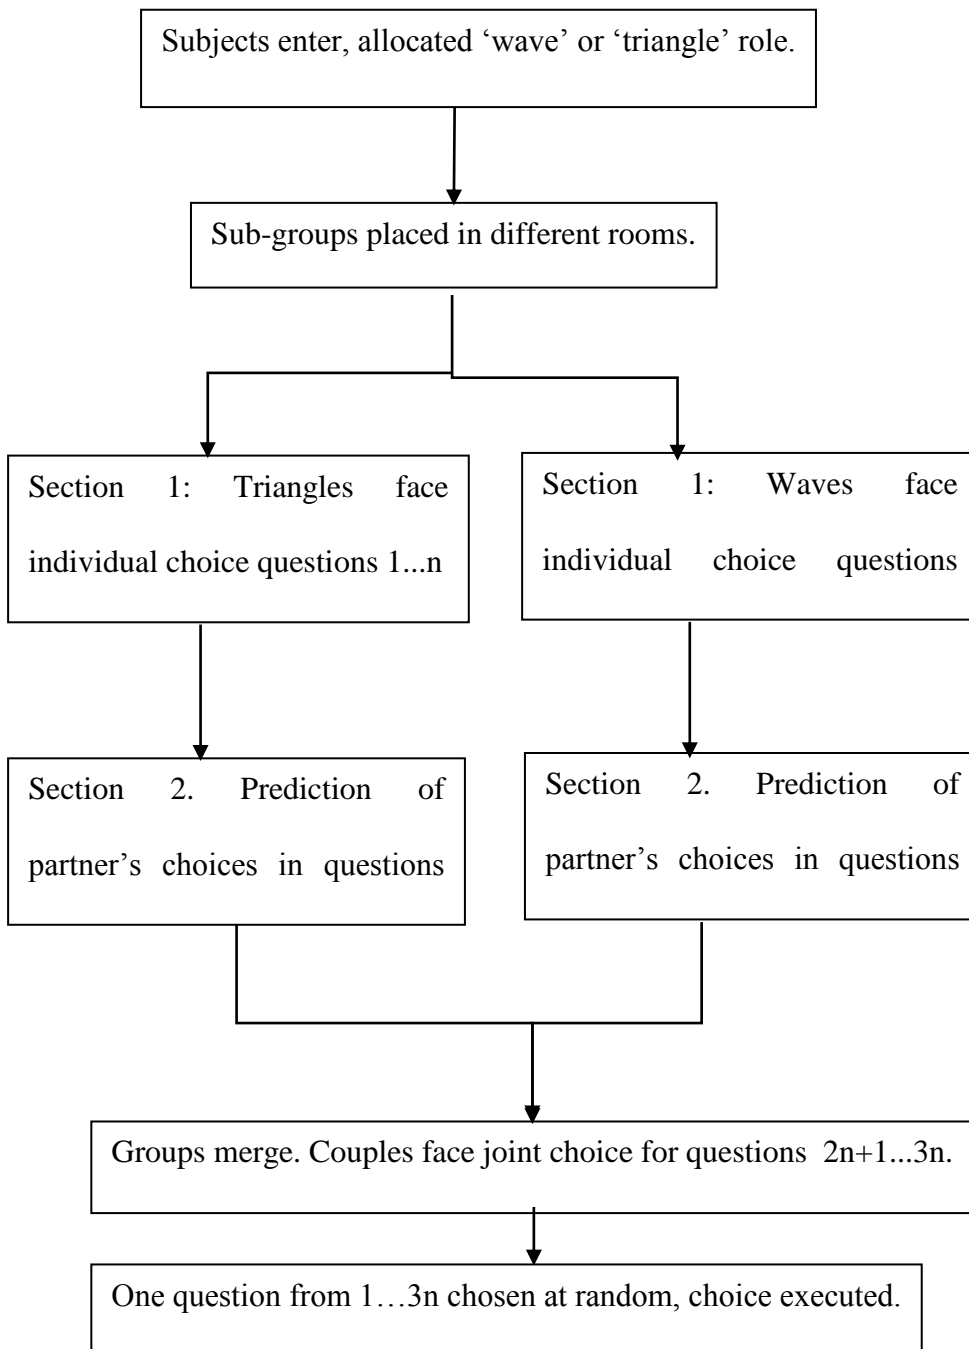


Figure 3. Experimental Procedure.

Appendix A. The tasks.

	Task type	Lottery 1				Lottery 2			
		Triangle		Wave		Triangle		Wave	
		£20	£40	£20	£40	£20	£40	£20	£40
Group 1	1 (T, W, J)	1-100	-	-	-	21-70	71-100	-	-
	2 (T, W, J)	21-100	-	1-20	-	21-100	-	71-100	-
	3 (W)	-	-	1-100	-	-	-	21-70	71-100
Group 2	4 (T, W, J)	1-100	-	-	-	-	-	-	41-100
	5 (W, J)	1-50	-	51-100	-	41-100	-	41-100	-
	6 (T)	1-100	-	-	-	-	41-100	-	-
Group 3	7 (T, W, J)	-	-	51-100	-	-	-	-	71-100
	8 (T)	51-100	-	-	-	-	71-100	-	-
Group 4	9 (T, W, J)	1-100	-	1-100	-	1-100	-	21-70	71-100
Group 5	10 (T, W, J)	-	-	1-100	-	-	1-70	-	-
Group 6	11 (T, J)	51-100	-	1-50	-	1-100	-	-	71-100
	12 (W)	31-100	-	1-30	-	-	71-100	1-100	-
Other	13 (T, W, J)	1-50	-	51-100	-	1-50	51-100	51-100	-
	14 (W)	1-70	-	21-70	71-100	71-100	-	1-70	-
	15 (J)	-	71-100	1-70	-	1-60	61-100	-	-
	16 (J)	-	-	21-70	71-100	1-40	-	-	71-100

Note: in this table the numbers shows the ranges of disc values for which the corresponding payoffs were awarded. To save space, we omit the numbers for states of the world where the payoff was zero. We also omit those tasks which are not relevant for this paper.