

Bateman, Ian; Munro, Alistair

Working Paper

Testing economic models of the household: An experiment

CSERGE Working Paper EDM, No. 04-01

Provided in Cooperation with:

The Centre for Social and Economic Research on the Global Environment (CSERGE), University of East Anglia

Suggested Citation: Bateman, Ian; Munro, Alistair (2004) : Testing economic models of the household: An experiment, CSERGE Working Paper EDM, No. 04-01, University of East Anglia, The Centre for Social and Economic Research on the Global Environment (CSERGE), Norwich

This Version is available at:

<https://hdl.handle.net/10419/80258>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

**TESTING ECONOMIC MODELS
OF THE HOUSEHOLD:
AN EXPERIMENT**

by

Ian Bateman and Alistair Munro

CSERGE Working Paper EDM 04-01

**TESTING ECONOMIC MODELS
OF THE HOUSEHOLD:
AN EXPERIMENT**

by

Ian Bateman^{1,2} and Alistair Munro³

**¹The Centre for Social and Economic Research
on the Global Environment**

**²School of Environmental Sciences
University of East Anglia
Norwich, Norfolk, UK**

and

**³School of Economics
University of Queensland
QLD 4072, Queensland, Australia**

Contact details:

**Alastair Munro: email - alistair.munro@uq.edu.au
Tel: 0044(0)1603 593404, Fax: 0044(0)1603 456259**

Acknowledgements:

The support of the Economic and Social Research Council (ESRC) is gratefully acknowledged. This work was part of the interdisciplinary research programme of the ESRC Research Centre for Social and Economic Research on the Global Environment (CSERGE).

ISSN 0967-8875

Abstract

We present the results of an experiment involving established couples, which uses choices between lotteries to test some economic models of household decision-making. Subjects make choices individually and jointly and are asked to make predictions about their partner's choices. Income pooling is not rejected in joint choice but has less explanatory power in individual choice. Many joint choices do not satisfy the Pareto principle. Overall, couples are more risk averse when making choices jointly compared to individual choice. Gender is not a direct determinant of power in joint choices, but female economic dependence significantly reduces women's decisiveness in joint choice.

Keywords: Household, economics, decision-making, experiment.

1. INTRODUCTION

Although there is a large body of experiments on individual choice, there is to date little formal experimental work that attempts to distinguish between economic models of household behaviour. Within psychology there is a tradition of work that attempts to find descriptive models of family decision-making, usually on the basis of data obtained from expenditure diaries, questionnaires and stated preference methods (see Eva Corfman and Daniel Lehman, 1987). However, the relationship between the results of this literature and the formal models of economists is at best unclear. This paper presents the results of an experiment designed for two purposes: first, to test properties of some of the main competing models of the household, including the Becker or unitary model and the Paretian or collective model and secondly to generate data for new or alternative theories. To do this we recruit established couples¹ and face them, individually and jointly with a sequence of binary choices over monetary lotteries which involve payoffs for themselves and for their partners. By comparing individual and collective choice we can test the Pareto principle. By manipulating which individual receives the payoffs while maintaining the household payoffs we test the principle of income pooling within couples.

Within economics the most commonly encountered formal models of household decision-making are the *unitary*, the *Pareto-efficient* or *cooperative* model and the *non-cooperative* model. In the unitary approach the household is modelled as a single agent with a single set of preferences - either 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. For instance, in Gary S. Becker's 'rotten kid theorem' models, 1974, assumptions about individual preferences and behaviour imply that all members of the household act to maximise its income.² The key feature of the co-operative approach, (see for instance, Marilyn Manser and Murray Brown, 1980, Martin Browning and Pierre-Andre Chiappori et al, 1998), is the assumption of Pareto efficiency in the intra-household allocation of resources.³ This assumption is relaxed in the non-cooperative class of models (e.g. Zhiqi Chen and Frances Woolley, 2001). Instead the Nash equilibrium notion is employed. Individuals make their contributions to household goods and choose their consumption patterns non-cooperatively. Hybrid models such as that proposed by Shelly Lundberg and Robert Pollak, 1993, have a cooperative solution coupled with a non-cooperative model of the threat point.

In recent years there have been a number of significant tests of household models using market data (see Cheryl Doss, 1996 and Harold Alderman et al, 1995 for surveys). Income pooling implies that, for given commodity and factor prices, patterns of expenditure should be independent of who receives the income in the household. In an important paper, Lundberg et al, 1997, use the natural experiment provided by changes in the payment system for the UK's child benefit to show that the identity of the income recipient does affect spending patterns on goods such as men's and women's clothing. Meanwhile, Shelley Phipps and Peter Burton, 1998, and Browning and Chiappori, 1998 use Canadian family income and expenditure data to show the systematic influence of the source of income on expenditure shares for many categories of consumer goods. One of the few experimental papers on the issue of household behaviour is an enterprising study by Elizabeth Peters et al, forthcoming. They employ family groups to test the 'rotten kid theorem' in a common pool game. Compared to the contributions

¹ By established, we mean living together and in a relationship of at least one year's standing.

² As pointed out by Theodore Bergstrom, 1989, the rotten kid theorem may fail to hold when the number of commodities is two or more.

³ When bargaining takes place in a risky environment, the co-operative models can be divided into two classes, depending on the feasibility of binding agreements over state-contingent intra-household transfers. If such transfers are possible, then the outcome is ex-post Paretian efficient; if they are not possible then only ex ante efficiency is possible.

made within groups composed of strangers they find that contributions made within family groups are typically higher and the reductions in contributions over time are much weaker. However, there is still some free-riding, particularly by children.

Tests of the co-operative model yield conflicting results. When published consumption data from developed nations is used (e.g. Browning and Chiappori, 1998) the unitary model is rejected in favour of the co-operative alternative. However, when data on production and factor supply is also used, typically from developing nations, both the unitary and collective models are often rejected. For instance, Christine Jones, 1983, rejects efficiency in household allocations in the Cameroon, Christopher Udry, 1996 finds differing rates of marginal productivity and marginal rates of substitution in data in farm level data in Burkina Faso and Udry, and Esther Duflo, 2001, also reject Pareto efficiency for farm households in the Côte d'Ivoire.

These differences in results may be the result of different economic contexts or may reflect the quality of available data. For instance, in many published data sets on family expenditure, variation in factor prices (particularly wage rates) is one potential but unmeasured source of variation in incomes. Even in Lundberg et al, 1997, which provides the cleanest test of income pooling in market data, the volatile nature of spending on clothing in the UK means that the authors have to average spending across several years in order to make a sensible evaluation of their hypothesis. Consequently, there is some scope for more direct tests of the predictions of household models.

One difficulty of testing economic theories of the household, is that, they differ not only in the equilibrium notion involved, but also in the level of detail given about the process which leads to the proposed outcome. So, it is not always clear how to apply some of the theories or to establish exactly what constitutes a fair test of their assumptions. Should a theory, for instance, apply only to the choices made by the household collectively, or should it apply equally to the choices made by individuals within that household? Consider, for instance, an individual who is a member of a household which makes collective choices according to $\mu(m_1, m_2)u^1(x_1) + (1-\mu(m_1, m_2))u^2(x_2)$, where u^i , $i=1,2$ is the riskless utility function for the i th individual, m_i is income received by the i th individual, x_i is the consumption of the i th individual and μ represents the weight attached to individual 1 in the collective decision. Suppose individual 1 must make a choice between (10, 2), (8, 12) and (2, 20) where (a,b) means £a paid to the individual and £b to their partner. They might choose according to $\mu u^1 + (1-\mu)u^2$ (the *collective* individual) or they might choose according to $u^1(m_1)$ (*myopic*) or they could choose according to u^1 knowing that any eventual allocation of resources will be according to $\mu u^1 + (1-\mu)u^2$ (*Stackleberg*). In the second case they will choose (10,2), but in either of the other cases they may choose one of the other two options, depending on μ and its relation to the values of m_i .⁴

Given this ambiguity over the exact predictions of the models, we therefore concentrate on testing some key properties that are features of a number of common variants of the main models. It is these properties we discuss next.

⁴ To be more specific, suppose that $u^i = \log(x^i)$, and $\mu = (m^1/(1+m^1+m^2))^{0.5}$, then the collective choice will be (2,20) and the Stackleberg choice will be (8,12).

2. PROPERTIES OF HOUSEHOLD MODELS

Consider a two-person household (or an n person household in which only the preferences of two agents matter directly for decisions). Index agents by $i=1,2$ and states of the world by $s=1,\dots,S$. Denote a lottery, L , where agent i receives payment m_{is} in state of the world s , by the vector $(m_{i1},\dots,m_{iS},m_{21},\dots,m_{2n})$ and let $L^i = (m_{i1},\dots,m_{iS})$, $i=1,2$. Say that two lotteries, L^A, L^B are Income Pooling Equivalent (IPE) if, for all $s=1,\dots,S$, $m_{1s}^A + m_{2s}^A = m_{1s}^B + m_{2s}^B$. Let \succeq be the weak preference relationship for the household (with strict preference, \succ , and indifference, \sim , constructed in the usual manner) and let \succeq_i represent the weak preference relationship for individual $i = 1,2$.

Four potential properties of preferences are:

Dominance (D). If for all i, s , $m_{is}^A \geq m_{is}^B$ then, $L^A \succeq L^B$ and if also $\exists s, i$, such that $m_{is}^A > m_{is}^B$ then, $L^A \succ L^B$.

Income Pooling (IP). If L^A, L^B are IPE then for all L^C , $L^A \succeq L^C \leftrightarrow L^B \succeq L^C$.

We shall call individuals income poolers if their choices satisfy IP. We shall also refer to lottery A as IP-safer (riskier) than B , when from an income pooler's perspective, A is safer (riskier) than B in the sense of Michael Rothschild and Joseph Stiglitz, 1970, meaning a lottery which has no risk when viewed by an income pooler.

Interpersonal Separability (S). For all, L^A, L^B , such that for one i $L^{Ai} = L^{Bi}$ and L^C, L^D such that $L^{Ci} = L^{Di}$ and, $j \neq i$, $L^{Cj} = L^{Aj}$ and $L^{Dj} = L^{Bj}$, $L^A \succeq L^B \Rightarrow L^C \succeq L^D$.

Pareto Dominance (P). $L^A \succeq_1 L^B$ and $L^A \succeq_2 L^B \Rightarrow L^A \succeq L^B$. $L^A \succ_1 L^B$ and $L^A \succeq_2 L^B$ or $L^A \succeq_1 L^B$ and $L^A \succ_2 L^B \Rightarrow L^A \succeq L^B$

The fourth of these properties is a relationship between individual and household preference. Although the other three properties have been stated for the household, with minor modification they can also hold at the level of the individual.

Dominance is a property shared by all major models of the household. However, it might not hold if agents are highly sensitive to issues of equity. Dictator and ultimatum games have revealed that strangers and acquaintances can sometimes reject dominating bundles, instead choosing dominated but fairer allocations (Werner Guth and Tietz, 1990). Given high levels of other-regarding behaviour within households, it is possible that similar choices of dominated bundles might occur.

Income pooling is usually associated with the unitary model, in which context both individuals and the household choose in accordance with IP and make the same choices. However, IP is also an implication of collective choice in the non-cooperative model provided both agents make strictly positive contributions to the household public goods and such contributions are made ex post - i.e. after the state of the world is known. In equilibrium, a transfer of Δm from one individual to the other in state s , leads to the receiving individual increasing his or her contribution to the household goods by Δm , while the other agent decreases his or her contribution by the same amount. For both agents, the net effect on consumption is zero for household and private goods, (Peter Warr, 1983).

Though IP is not generally an implication of the Paretian model it is compatible with it under specific conditions. Suppose for instance 1) that any state-contingent transfers between the agents are possible 2) that the household chooses such transfers and other variables under its

control to maximize $\mu E(u^1) + (1-\mu)E(u^2)$ where $E(\cdot)$ is the expectations operator and 3) μ is independent of the pattern of individual income.⁵ Under these assumptions, any external transfer of Δm from one individual to the other can be undone by adjusting the pattern of intra-household transfers. Hence, IP will hold for such a household. However, IP will not hold when intra-household transfers are constrained to be ex ante.

The third property is one of additivity. It implies that the household or individual will be indifferent about the correlation between L^{A1} and L^{A2} . So, if for instance, the household maximizes $\mu E(u^1) + (1-\mu)E(u^2)$, ex post transfers are not feasible, and μ is insensitive to changes in the pattern of household income then S will hold. However, when ex post transfers are possible, the set of feasible transfers may be changed by the pattern of correlation between L^{A1} and L^{A2} . It follows that in ex post efficient models of the household or in non-cooperative models where expenditure allocations are made ex post, S will not hold in general.

The final property is a key prediction of all Paretian models. It implies that if two individuals both prefer lottery A over B, then acting jointly as a household, they will not prefer B to A. One thing to note is that none of these properties make the assumption that either individual preferences or the preferences of the household satisfy the assumptions of expected utility theory (EUT). EUT has a long history of experimental scrutiny, most of which is hostile to the assumptions of the model. As a result we design our experiment without assuming EUT.

⁵ Typically in Paretian models μ is not independent of the pattern of income. Indeed, it is a central purpose of many such models to explain μ as a function of income and other determinants of bargaining power. However, for small changes in income shares we might reasonably expect μ to be approximately constant.

3. EXPERIMENTAL DESIGN

The data examined in this paper were generated by two very similar experiments, A and B which shared the format summarised in Figure 1. In the first and second sections of each experiment, individuals were separated from their partners and assigned the roles of 'triangle' or 'wave'; in the third section they rejoined their partner. In the first section each individual faced n tasks, each comprising a choice between a pair of lotteries; in the second section each individual was asked to make predictions about the n choices their partner had made in their section 1 and in the third section each couple made choices between n pairs of lotteries. The n tasks were not all repeated in each, though there was some overlap (see Appendix A). The experimenters used a script and subjects received summary instructions separately for each part of the experiment. Appendix B shows the written instructions given to subjects in the first section of the experiment.

In experiment A, $n=10$ and in experiment B, $n=12$. Twenty-four questions were common to experiments A and B - we added six and replaced six. The changes were made for two main reasons: first, it became clear that the experiment could accommodate more questions in the self-imposed one hour time slot⁶; secondly, some hypotheses were quickly rejected by the accumulating data and we took the opportunity to drop the majority of the relevant tasks.

To determine payoffs a question number from 1 to $3n$ was selected at random. If the number was between 1 and n , the triangle partner played his or her choice for that question.⁷ Meanwhile the wave partner received fifty pence for each correct prediction she or he had made.⁸ If the question number was between $n+1$ and $2n$, the wave partner played his or her choice for that question. Meanwhile the triangle partner received fifty pence for each correct prediction she or he had made. Finally, if the question number was between $2n+1$ and $3n$, then the couple played their choice for that question and no money was paid for predictions.

This scheme is incentive compatible if individuals are selfish in their choices and binding agreement on ex-post trade is not possible. However, it is not unreasonable to suppose some degree of other regarding behaviour in members of a couple.

Theoretically, an altruistic individual might view the first two sections of the experiment as an exercise in co-ordination. They could make their choices so as to maximize the predictive success of their partner. Three aspects of the experiment were designed to guard against this possibility. First, subjects were not informed about the detailed nature of the experiment until after they were separated from their partners. In the recruitment literature we merely stated that the experiments aimed 'to help us understand how couples make decisions'. In the brief welcoming statement we explained only that the experiment was about choice. Secondly the prediction section always came after the choice section and subjects were not informed about the type of tasks in section 2 until they had all completed section 1. So, subjects had no reason to anticipate that they should tailor their answers in section 1 so as to raise the possible payoffs of their partner. Finally, the payments for prediction were much lower than those associated with choice: the maximum payoff from the predictions was $n/2$ pounds whereas payoffs from each individual in the choice section were up to £40 per person. The expected values of the two options on offer always differed by more than fifty pence. So, as a benchmark example, a risk

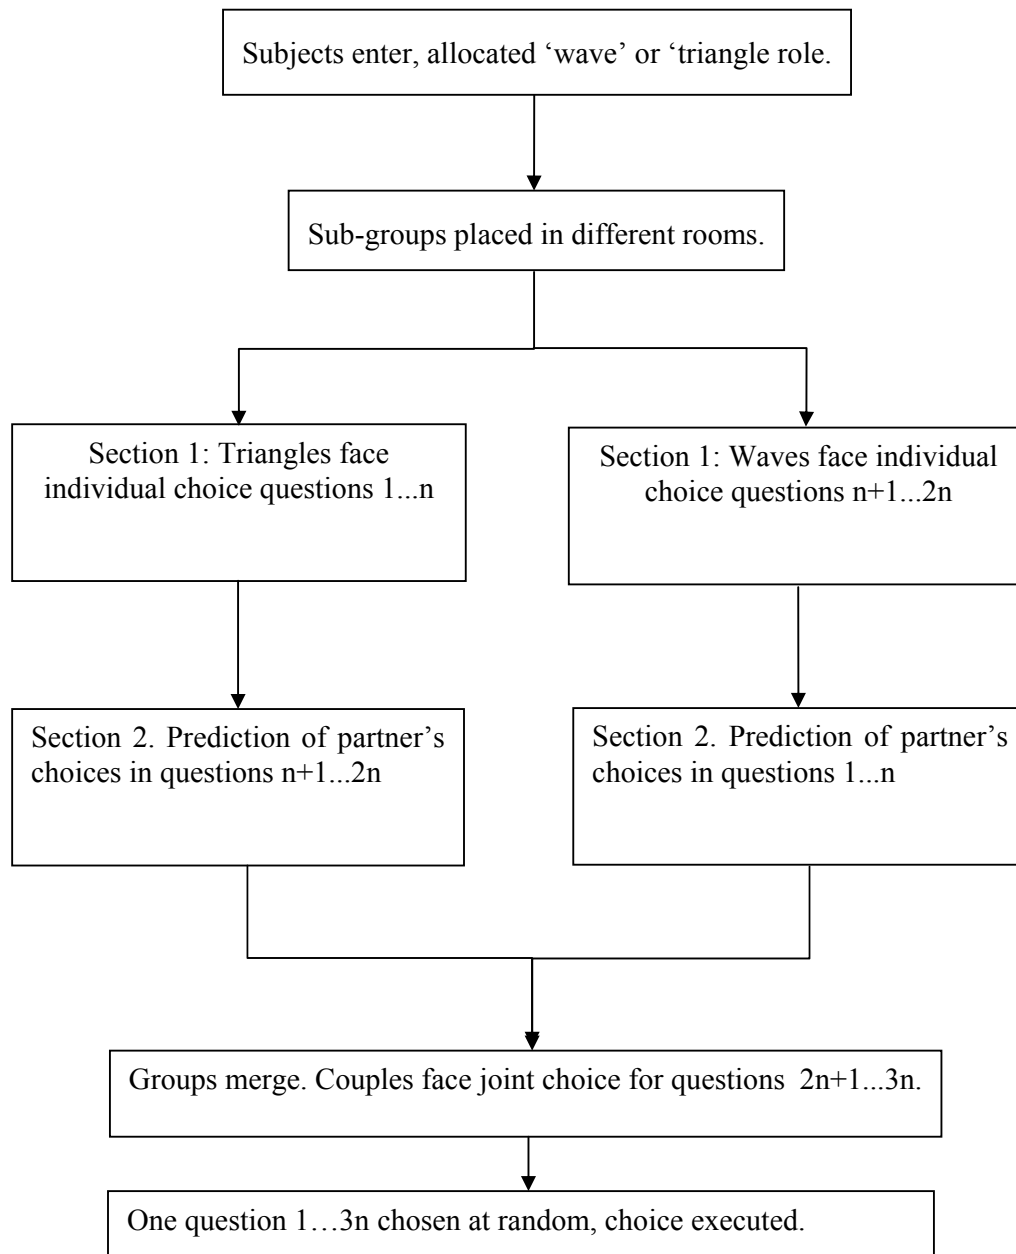
⁶ Many of the sessions were held at lunchtime or after work. We were conscious of the time constraints on our subjects, most of whom were in paid employment and/or had children, so we designed the experiment to fit into a one hour slot.

⁷ Robin Cubitt et al 1998 provide evidence that even when subjects are not EUT maximizers, the choices made in experiments with random lottery incentives are not sensitive to the value of n .

⁸ This means that the wave partners saw questions in the order $n+1$ to $2n$ (choice) and then 1 to n (prediction).

neutral income pooling agent would never gain from switching choices in order to improve the predictive success of their partner.

Figure 1: Experimental Procedure



Subjects were not informed of their partner's answers in section 1 of the experiment; nor were they given information about the accuracy of their partner's predictions.⁹ In fact, we took considerable trouble to preserve the confidentiality of their answers from sections 1 and 2. The

⁹ Given our incentive system, if one member of a partnership got all his or predictions correct (or none correct), then provided she or he had perfect recall of the questions, that subject could know a partner's choices. However, no participant raised this possibility with us during the conduct of the experiment and no individual achieved perfection in their predictions (or scored zero).

theory of the household is rarely explicit on whether individual members of a household are privy to the patterns of consumption and income of their partners. Empirically though, there is plenty of evidence of asymmetric information in spending behaviour. For instance, in her survey of spending habits in a sample of UK households, Jan Pahl, 1990 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 Judith Treas, 1993 for similar US evidence and Frances Woolley, 2000, for Canada). Thus it is reasonable to expect theories to be robust in the face of such possibilities. However, a fundamental reason for adopting confidentiality is to avoid producing incentives for individuals to misrepresent their preferences over lotteries in order to avoid recrimination or procure approval from their partners. Many economic theories of the household relate individual preferences over bundles of goods to patterns of collective choice. Thus, in order to test such theories we typically require data on individual preferences over bundles of goods. Revealing choices to partners might instead generate data on preferences over actions (e.g. 'pleasing my partner'). Such preferences may be an important source of household behaviour, but they are not typically the objects of preference in economic theories of the household. As a result, we chose to protect the confidentiality of individual responses.

As subjects entered the room each member of the pair was given one from a pair of cards at random. One card showed a 'wave' design and one showed a 'triangle'. These cards were used to indicate roles in the experiment that followed. After briefly introducing the experimenters, each subject was given a large sealed envelope marked with either a wave or a triangle and instructed not to open it. Then, one half of the subjects were led to a separate room for the first two sections of the experiments. Once the subjects were separated into their respective groups, the subjects were told to take out of the envelopes the colour-coded booklet containing summary instructions and answer sheets for the first section of the experiment.

The instructions for the first section explained the nature of the lotteries that were the focus of the experiment. Subjects were also told that at the start of the third section, each couple would choose a small envelope which would contain a number to be revealed at the end of the experiment. If the number matched one of the question numbers for section 1, then they would play out their chosen lottery from that question. Subjects were not told at that stage what would happen if any other number was in the small envelope.

Figure 2 illustrates a typical question from the question. As can be seen, each depiction of a lottery was composed of three elements: ranges of numbers along the top; payoffs for the subject below and then payoffs for their partner.¹⁰ The subjects were informed that the numbers along the top corresponded to numbers discs in a bag of one hundred discs shown to them by the experimenters. They were told that if a question number from section 1 was in their small envelopes then they would choose a chip from the bag and this would yield the corresponding payoffs for them and their partner.

At the end of the instructions for the first section we placed two questions designed to test understanding. Once all subjects had answered these questions satisfactorily, which nearly all did at the first attempt, they were instructed to answer the choice questions in their own time.

When all subjects in the group had completed all the questions for section 1, they were led through the instructions for section two, including those concerning incentives. At the end of the answer sheets for the second section was a short questionnaire with questions on age, gender, number of children and domestic financial arrangements. Once all subjects in both groups had completed this as well as their prediction questions they rejoined their partners for

¹⁰ For the jointly answered questions, the triangle partner's payoffs were always shown first. We found no evidence (see end of the next section) that this order gave triangle partners more or less influence in the joint decisions.

the final section of the experiment.

At this stage the experimenters took away the envelopes marked with triangles and waves and gave each couple two envelopes - one small and one large. The large one contained the instructions and answer sheet for the third section. Meanwhile, the couples were invited to choose a small envelope from shuffled pile placed in front of them and told not to open it until instructed. No prompts were given as to which member of the couple should make the choice. Inside the small envelope was the number that would determine which question would be 'for real' at the end of the experiment. Lottery ticket numbers from 1-3n had been previously allocated at random to the small envelopes.

Before the subjects were told to open their large envelopes, they were given more explicit details of how the payout procedures would operate at the end of the experiment. If the question number in the small envelope was from section 3, then one of the experimenters would oversee the choosing of a disc from the bag and payouts in situ. However, if the number was from either section 1 or 2, then in order to preserve confidentiality the partners would be dealt with separately, beginning with the chooser. He or she would be asked to go to a table at the far end of the room (or in another room for some sessions) where they would be reminded of their choice and a disc would be drawn from the bag. Following this he or she would receive any payments owing to them in an envelope and then asked to leave the room, while their partner was dealt with. Again any money owing to the partner was placed in an envelope to protect confidentiality.

The subjects were lead through the instructions for the final section and then asked to complete their final set of answers. Once all subjects had completed the tasks, the experimenters began the task of opening the small envelopes and making the appropriate payoffs.

The two experiments were carried out over the period from December 2002 to March 2003, with a brief pause in between. After a pilot session,¹¹ subjects were recruited from and around the city of Norwich via email, through community groups and using posters. Session sizes varied from 2 couples 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. The majority of the sessions were held at lunchtimes, but approximately one third of subjects did the experiment in the evening. 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 evidence¹² of their relationship and made random checks.

¹¹ Following which we fine-tuned the script and probabilities, but made no structural changes.

¹² Most couples had been together for much more than one year (see Table 1). Evidence brought to the sessions included photographs, bills to the same address and, in three cases, children.

Figure 2: A Typical Question

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

☐

4. RESULTS

In all 42 couples produced the information for dataset A and 34 couples produced information for data set B. Average payoffs were just under £17 per individual, which is more than twice the median hourly post-tax wage for a UK adult. Standard statistical tests indicated no significant differences between the data for the questions common to both sub-samples. Consequently, the data is routinely pooled in what follows.

Table 1 provides some background information about the individuals in the experiment, based on the answers to the socio-economic questions which subjects faced at the end of section 2. The data is presented on an individual basis (n=152) because that was the way it was collected. As can be seen, the range of ages and years together was quite diverse. All the couples in our sample were heterosexual with seventy-three percent married to their current partner. The distribution of children per couple was bimodal with peaks at zero and two. For the financial question, if subjects ticked the 'other' box they were asked to write in their description of their domestic financial arrangements. The most common such arrangements were of two kinds: one involved three accounts - one for regular household bills and separate accounts for personal expenditure. The other common arrangement was a decentralised system for routine expenditures, with one partner responsible for some household bills while the other partner paid for other items. This was often coupled with joint decision-making for major or idiosyncratic events. For instance, one individual wrote, 'we have a joint account to cover all household and shared expenses. We retain some money each to spend as we wish.' Another stated, 'accounts separate- each paying for different bills. i.e. one pays for mortgage and other car insurance and utilities.'

On a regular basis, partners gave different answers to the same question, particularly when asked about years living together and financial arrangements. Although no couple differed by more than two years, 33% stated different lengths for their relationship. Meanwhile, 43% of couples differed in their assessment of their financial arrangements. Overall, just 39% of couples gave the same response to these two questions.

Table 1: Sample Characteristics

	Mean	Standard deviation	Range
Children	1.1	1.19	0-4
Age	37.3	10.74	22-70
Years living together	11.1	10.59	1-46
Best description of financial arrangements:			
I look after the household money except my partner's personal spending money;			15
My partner looks after the household's money except my personal spending money			12
I am given a housekeeping allowance. My partner looks after the rest of the money;			
My partner is given a housekeeping allowance. I look after the rest of the money			
We share and manage our finances jointly			87
We keep our finances completely separate			13
Some other arrangement			21
Note: Two individuals gave multiple answers to this question. They are omitted.			

4.1 Main Results

We now turn to the central results from the experiment. In what follows, the tasks are labelled. The letter T indicates a task faced by triangle subjects, W stands for tasks faced by wave subjects and J indicates tasks put to couples in the final part of the experiment. Since there were different tasks in experiments A and B and the question order was different, the numbers attached to tasks do not correspond to the order in which any subjects saw them, but are there for reference purposes. The details of each task can be found in the Appendix.

Result 1. Subjects rarely choose stochastically dominated options.

When triangle subjects faced the choice (T13) depicted in figure 1, 4 out of 76 (5%) chose the dominated option. When couples faced the same choice (J11), 2 out of 76 (3%) chose the dominated option. When wave subjects faced the same task (W13), 8 out of 76 (10%) chose the dominated option. In the fourth test of dominance (W14), 2 out of 42 wave subjects chose the dominated option.

Result 2. The combination of P and S is rejected.

Table 2 shows the results for all of the tasks for which the combination of the Pareto principle and Interpersonal Separability makes a clear prediction of one option, but where that alternative does not stochastically dominate the other. In four of these tasks, for one individual options A and B are identical. In the fourth and last task, when viewed individually one option strictly dominates the other option for both individuals. However, as can be seen, in none of the tasks does the proportion who choose the option predicted by the theory rise above 50%. This is true whether the choice is made by individuals or by couples.

Table 2: Within task tests of P and S

Task	J12	T2	W2	J2	W9	J8
Instances predicted option actually chosen	24	16	10	10	10	4
N	76	34	34	34	34	34
Percentage	31.6	47.1	29.4	29.4	29.4	11.8
Chi-squared	493**	124**	236**	236**	236**	383**

** statistically significant at 1 per cent level, d.f. = 1

The chi-squared values in the table are generated in the following way. To test formally whether P and S are jointly rejected, we need to make some assumptions about the stochastic process underlying agent's choices in the experiment. If agents are assumed to choose without error then any deviation from the 100% predicted by the theory for these tasks would imply rejection of the maintained hypothesis. This seems too extreme, so in keeping with other models of choice in an experimental setting, we posit some degree of agent error. Specifically, we suppose that the choice of stochastically dominated options represents only the results of a 'tremble' in which subjects mean to choose the dominant option but accidentally pick the dominated lottery. As a benchmark figure we use the 6 per cent average figure from the three identical tasks discussed above in the context of stochastic domination.¹³ As can be seen, using this figure, for each task the null hypothesis is rejected at any widely recognised level of

¹³ Arguably we should use a lower figure for couples, on the grounds that 'two heads are better than one' when it comes to spotting errors. For simplicity we stick to the same figure for both individuals and couples.

statistical significance. Moreover this conclusion is not sensitive to the degree of error posited. All of these tasks share one key feature: the alternative, apparently dominated option is risk-free when viewed by individual agents and couples who pool their income. In the vast majority of cases, it is this 'safer' option which is chosen by the subjects. So, P and S are not only decisively rejected by the data as a pair of joint assumptions, it seems that, as an assumption, IP has superior explanatory power. It is this issue we consider next.

Result 3. On the basis of within task tests, IP is not rejected for joint choice, but has less explanatory power for individual choice.

Table 3 shows the results of six questions where, for an income pooler, one option strictly stochastically dominates the other. To a large degree subjects choose according to the predictions of the theory. In the three jointly chosen options, the income pooling option is almost always chosen. In fact, the pattern of answers is almost exactly that which would be predicted with a tremble of 6%. In the case of individuals, it is still the case that the overwhelming majority of choices are compatible with income pooling, but there is now a substantial number who do not choose according to the theory. As a result, with a tremble of 6%, income pooling is rejected for one of the four individual choice tasks in the table and as the final row of the table shows, it is rejected when we pool all the relevant individual choice questions.

Table 3: Within task tests of IP

Task.	T11	T12	W11	W14	J10	J13	J14
Instances predicted	59	37	69	40	71	39	40
option chosen							
N	76	42	76	42	76	42	42
Percentage	77.6	88.1	90.8	95.2	93.4	92.9	95.2
Chi-squared (d.f. = 1)	34.0**	2.42	1.24	0.14	0.03	0.07	0.14
Chi-squared for joint test across:	Individual choice: 21.3**				Joint choice: 0.02		
** indicates statistically significant at 1% level							

The difference in results for the joint and individual choice cannot be simply attributed to differences in tasks. The first individual choice shown in the table and the first joint choice are actually identical tasks, but nevertheless yield very different patterns of behaviour. One possible reason for the difference is that the minority option in these two tasks is extremely equal: it gives each partner a 50% chance of winning £20. To the extent that subjects, when acting as individuals, are motivated by considerations of family equality, then this would account for the significant number of individuals who opted for it. In support of this interpretation it is worth noting that the fourth individual choice task in table 3 is identical to the first from the viewpoint of the income pooler, but the option chosen by the minority does not, in this case, show the same pattern of equality and did not yield the same proportion of choice.¹⁴

Result 4. On the basis of between task tests, IP is not rejected for joint choice, but has less explanatory power for individual choice.

¹⁴ It is also worth noting that the behaviour of many of the choosing agents in the case of the first individual choice task was not anticipated by their partners, who overwhelmingly (72 subjects out of 76) predicted that the choosers would act in line with income pooling.

Having considered within task tests of income pooling, we now turn our attention to seven groups of tasks within which all the tasks are IP equivalent. Table 4 provides some basic results on the seven groups of tasks, which we explore more thoroughly below. IP agents should choose consistently within a group and so the proportion choosing the IP equivalent options should be the same across all tasks within any given group. The significance tests are indicative and should be interpreted with care. They show all proportions which are significantly different from at least one other task within the group. However, note that if there are n tasks within a group, then there are $n(n-1)/2$ possible tests within that group and so, even if the null hypothesis of IP is true, we would expect some instances of significance.¹⁵

For the three groups (1, 2 and 3) where it is possible, we test the hypothesis that the proportion jointly choosing the IPE option is the same. In each case the null hypothesis is not rejected. The *lowest* p value for the chi-test is 0.38. Also for each of the seven groups we test the hypothesis that the proportion individually choosing the IPE option is the same. In this case the results are mixed: for group 7, the null hypothesis is rejected ($p < 0.001$), but it is not rejected at any standard significance level for any of the groups. For each of the seven groups we also test the hypothesis that the proportion choosing the IPE option is the same across both joint and individual choice. This hypothesis is rejected for groups 2 and 7 ($p < 0.001$), but not rejected for the other groups.

Of course, for these tasks, IP does not just predict equal proportions in the sample; it also predicts consistency within individuals. Figure 3 summarises the information from the six pairs of IPE tasks where $N=76$ - five involving individual choice and one involving joint choice. In each case the lower bar shows the number of individuals who would be listed as consistent if the choice probabilities across tasks were independent. As can be seen, in the main it is substantially lower than the actual figure and when we use a chi-squared test, we reject independence at the 1% level for all but the first and the last pairs shown in the figure.

The results so far suggest that IP is a reasonable assumption to make in the context of joint choice, but that for individuals, IPE tasks do not always yield equivalent choices. Moreover the choices made by individuals and couples do not always match. We now examine subsets of IPE tasks in order to understand how non-equivalence comes about.

Result 5. Couples are more risk averse in their choices than individuals.

The evidence for this claim is summarised in Figure 4 which shows the proportion choosing equivalent options for the seven sets of tasks where there is data on choice from both individuals and also from their joint choice as a couple. In other words, within each set the tasks are not just IPE, they are actually identical. The first two sets both come from Group 1 in Table 3; there is one set from each of Groups 2 to 6. With the exception of set 4, it is possible to rank the riskiness of the options from the viewpoint of an agent satisfying IP. For set 4 the riskiness of the options cannot be unequivocally ranked. For sets 1-3, 6 and 7 one option has no risk for the income pooler. For option 5, one option is safer.¹⁶

¹⁵ Standard ANOVA methods are inappropriate here given that the groups contain a mixture of within and between subject data.

Table 4: Cross-task tests of IP

Group 1	T1	T2	W1	W8	W2	J1	J2
N	76**	34	76**	76	34**	76**	34* *
% choosing equivalent option.	50	53	50	55	71	64	71
Group 2	T8	T3	W9	W3		J3	J8
N	76***	76***	34***	34***		76***	34* **
% choosing equivalent option.	63	78	71	76		93	88
Group 3	T4	T9	W4			J4	J9
N	76	76	76			76	76
% choosing equivalent option.	63	71	72			70	67
Group 4	T10	T5	W5			J5	
N	76	76**	34			76**	
% choosing equivalent option.	57	46	53			64	
Group 5	W6	W10	T6			J6	
N	76***	76***	34			34**	
% choosing equivalent option.	39	22	35			44	
Group 6	T7	W7				J7	
N	34	34**				34**	
% choosing equivalent option.	59	47				74	
Group7	T11	W11				J10	
N	76**	76**				34**	
% choosing equivalent option.	22	9				6	

All tests 2-tailed; ** = significant at 5% level; *** = significant at 1% level.

¹⁶ In all but one case (set 2) this is also the option which is safer when payoffs are viewed individually. In set 2 it is not possible to order unequivocally the riskiness of the options from an individual perspective.

Figure 3: Consistency with Income Pooling across six pairs of tasks.





Figure 4: Proportions choosing equivalent option.



For all the sets, except set 4, the bars in the figure show the proportions choosing the IP safer option. If the unitary model is true then each bar should be the same length. Clearly they are not. In fact there seems to be a systematic tendency for couples to choose the safer option more frequently than individuals. In all cases, with the notable exception of set 4, the bar is highest for the joint choice. On a one tailed test chi-squared test, the triangle proportions are significantly lower than the jointly-chosen proportions for sets 1 ($p=0.020$), 2 ($p=0.037$), 3 ($p < 0.001$), and 5 ($p=0.023$). On the equivalent test for wave versus jointly-chosen proportions, the significantly different sets are 1 ($p=0.020$), 3 ($p=0.005$), 5 ($p=0.044$) and 7 ($p=0.040$).

It is not clear why couples should tend to be more risk averse than individuals. Results 3 and 4 suggest that couples behave more like income poolers than individuals. Possibly therefore, the safety of one of the options is more salient when decisions are made jointly. Alternatively, perhaps safety is more attractive in joint decisions because choosing the safer option minimises the risk of recrimination once the state of the world is revealed. One possibility we can discount is that individuals routinely underestimate their partner's degree of risk aversion, which is then corrected when they are asked to make joint choices. The relevant prediction data for the seven sets in figure 4 shows no systematic tendency to over- or under- predict the degree of risk aversion. A final possibility lies in the possibilities for negotiation opened up by joint decision-making. When making individual choices, subjects have to make conjectures about ex-post trades, which might be false, and they are also denied the possibility of ex ante trade. However, given the opportunities for risk-sharing provided by negotiation, this kind of explanation would sit more easily with individuals making more risk averse choices compared to those made by couples.

With regard to IP there are some other comparisons that we can also examine, though the conclusions to be drawn are less clear. A lottery can be transformed along two dimensions in ways that leave the result IPE to the original. First, it can be made more or less equal. In an equal lottery, for each state of the world, each person receives the same monetary payoff. In the other extreme, for each state of the world the same person receives all the available income. The other dimension is individual risk. If a lottery is IP safe, then increasing the risk for one person means also increasing the risk for the other person. Conversely, if a lottery is IP risky, then decreases in risk for one person may imply increases in risk for the other person. Within the seven groups of tasks in table 4 there are various combinations of transformations represented. Particularly interesting are the two pairs (W1, W8) and (W6, W10) where the transformation is relatively simple. These two are examples where one task is the *mirror* of another, in the sense that in the transformed pair of lotteries, for each state of the world s partner 1 receives what partner 2 receives in the original lottery and vice versa. In these two cases, the effect of mirroring is to change the agent who directly faces the trade-off in risk. For the other agent both options are the same and both are completely risk free. For W1 the wave agents choose the risky option 38 times out of 76, when it is their partners who directly face the trade-off in risk. Conversely, in W8 wave subjects choose the option which is risky for themselves 34 times out of 76. This difference is not statistically significant ($p=0.26$, two-tailed paired-sample t-test). With the pair (W6, W10) the tendency is in the same direction: for W6, wave subjects choose the option which is risky for their partners 30 times out of 76, but in W10 they choose the option which is risky for themselves only 17 times out of 76. This difference is statistically significant at the 1% level ($p=0.01$, two-tailed paired-sample t-test). So, there appears to be some tendency for subjects to be more risk loving in choice when it is the partner who faces the trade-off between safety and risk.

The IPE tasks, T3, W3 and T8 present evidence from another set of simple transformations. For T8, triangle subjects face a choice between £20 for sure or a 60 per cent chance of £40 for themselves. Nothing accrues directly to their partner in either option. In T8, 63 per cent of subjects choose the safe option. In T3, triangle subjects must choose between £20 for sure for themselves and a 60 per cent chance of £40 for their partner. Here 78 per cent choose the £20. This difference ($p=0.03$ on a two-tailed paired-sample t-test of the null hypothesis that the

proportions are the same) might be attributed to subjects placing greater weight on their own payoffs than on their partners'. However, W3 is the mirror-image of T3 - wave subjects must choose between £20 for sure for their partner or a 60 percent chance of £40 of themselves. An explanation based on 'less than equal weight for partner' would predict that the proportion of wave subjects choosing the safe option would be lower than the proportion of triangle subjects making the same choice in T8. In fact 76 percent of wave subjects choose the safe option in W3 - almost exactly the same as that for T3. Possibly therefore, this switch towards the safe option is a heuristic brought on by the difficulty of making the simultaneous trade-offs between risk and safety and between self and partner.

Summing up: the rejection of separability showed that at least to some extent the individuals in our sample sum incomes by state of the world. However, especially for individuals choice is affected by who receives the income and who bears the risk.


Result 6. Revealed preferences do not provide evidence in support of the Paretian model for many couples.

According to P, if both individuals strictly prefer lottery A over B or one individual is indifferent and the other strictly prefers A, then as a couple they should not choose B. Figure 5 shows the pattern of agreement within couples for the seven sets of identical tasks in the experiment. For each set, the first bar shows the number of couples where individual choices differ. The second bar shows the number of cases where individual choices are congruent and are the same as the choice made jointly. Leaving aside the possibilities of indifference or error, if the Paretian model is correct, there should be no cases where the joint choice differs from the individual choices, when those individual choices are the same. However, as can be seen from the third series of bars, in each data set there are a number of couples whose joint choice does have this property. On average 23 per cent of couples with identical individual choices plump for a different option in joint choice.

Table 5: Testing Pareto: Cases where individuals choose the same, but couples do not

Set	1	2	3	4	5	6	7
No. cases where individuals agree	40	28	24	43	27	20	16
Of these, number where couples differ	9	4	2	7	8	8	6
Chi.	19.31	8.40	0.23	8.06	26.72	40.99	28.15
p value	1E-05	0.004	0.630	0.005	2E-07	2E-10	1E-07
Critical error rate	12.3	9	2.3	8.1	15.9	21.9	18.5

Figure 5: Agreement of preferences



One theoretical possibility is that in all the relevant cases both individuals are indifferent between options A and B, so that the joint choice does not represent a rejection of the Pareto principle. Alternatively, and perhaps more realistically, the pattern could represent the result of some random element in preferences.¹⁷ Table 5 provides chi-squared tests, based on a 6 per cent tremble as used above. For all but one data set, the null hypothesis is decisively rejected.

There are reasons to show some scepticism about the test in this context. First, if all choices are subject to a tremble of six percent, then some of the apparent agreements between individual choices used as the basis of the test will actually reflect previous errors. Depending on how disagreement on preferences are resolved this may raise or lower the expected rate at which joint choice should diverge from individual choices. Secondly, the error rate of 6 percent is based on data from tasks where one option stochastically dominates the other. Where choice is more evenly poised we might expect a greater degree of randomness. This is a common feature of random utility and random preference models. We do not have a repeated task in the data set which would provide a benchmark estimate of the random element of preference. However, in the case of joint choice, couples were faced with two tasks which were a mirror image of one another. In this pair 18 per cent of couples were inconsistent. If this figure is taken as the error rate then the null hypothesis would be rejected for only the last two data sets.

¹⁷ Another possibility is other regarding behaviour combined with an error in prediction which is then corrected once individuals are able to communicate with their partners. While we cannot rule this out, there is no supporting evidence for the hypothesis to be found in the prediction data.

Against the randomness interpretation, it is worth noting that there is a clear pattern of change between individual choice and joint choice. As was the case with figure 4, the reversals of preference are not random. For set 4, where neither option is definitively safer than the other, three couples switch in one direction and four in the other. For the remaining six sets, where in each case one option can be ranked as the safer option, twenty-eight of the switches are towards the safer option in joint choice and only nine are towards the risky option. In other words, as we saw above, when faced with the same tasks, couples seem to be more risk averse in their behaviour than individuals. In some cases this manifests itself in a complete switch from the riskier option chosen by both individuals to the safer option.

Result 7. Unanimity is not a feature of the majority of choices and few households are unanimous in all their choices.

By unanimous we mean that individual choices agree and also agree with the corresponding jointly-made choice. Out of 322 cases where we have data of this type, unanimity is present in 144 (44.7 per cent). For the 34 couples who face the seven complete sets of tasks discussed above, none are unanimous in all seven tasks - three are unanimous in six tasks. For the 42 couples who face just two complete sets, 5 are unanimous in both.

Result 8. Gender-related power is not per se a major determinant of joint choice, but pre-retirement age women without paid employment have significantly less influence in joint choices compared to women with paid employment.

A perennial issue is the degree to which household choices represent the wishes of one partner rather than the other. Psychological studies (see Corfman and Lehman, 1987, for instance) suggest that influence is not simply determined by gender but by a variety of other factors including perceived expertise and intensity of preference. The primary aim of our experiment was not to test for such factors, but nevertheless the data yields some interesting insights. Within the seven sets of identical tasks, there are 181 instances (out of a possible 364) where the individual choices differ from one another. Out of these, 86 joint choices reflect the individual choices made by the female partner (47.5%). This is not significantly below fifty percent. In three of the seven sets of tasks the percentage where the women are decisive is below fifty percent; in three sets the percentage is above fifty percent and in one case men and women are equally likely to be decisive. We conduct a probit analysis using as the dependent variable 'female decisive' - which takes the value 1 when the joint choice reflects the individual choice made by the female partner rather than the male. We pool the data across the questions and use question dummies. Being married, self-reported financial arrangements, length of time together, role in the experiment (triangle or wave), age and number of children have no significant impact on the probability that the female is decisive. None of the question dummies are significant except one. We investigate whether other measures of power have significant explanatory power. Specifically we construct a new variable of female financial dependence (findep) which takes the value 1 if the female member of the household states that she is a housewife or describes herself as having no occupation.¹⁸ Overall, females who were classified as financially dependent were decisive in one third of the relevant instances, compared to 79 out of 160 instances for other women. We also construct 'Agedep' which equals (findep) x (age of the triangle partner).¹⁹ We end up with the following equation (t values in parentheses):

¹⁸ For pensioner couples or where one person is retired, it is not clear who has external sources of income, so we do not code such households as showing female financial dependence. When we drop the ten cases involving retired females our estimation yields the same conclusions. There is one couple where the male describes himself as a househusband. Our results are also not sensitive to how we code this household.

¹⁹ We also construct a number of other variables related to the number of children, financial dependence, financial arrangements etc. None of these are significant, although in some specifications, a variable which

$$Y = -0.041 + 0.043\text{Agedep} - 2.47\text{Findep} + 0.75\text{Q29} \quad \text{Log-likelihood} = -120.$$

$$(-0.39) \quad (1.69) \quad (-2.05) \quad (1.98)$$

In this equation, Q29 is a dummy for one question where females were decisive in 10 out of 14 cases. Otherwise this task shows no obvious distinctive features. The equation suggests that female financial dependence is a significant factor in determining her lack of decisiveness in joint choice, but that the effect is attenuated by age, disappearing by age 57.

Result 9. Amongst individuals, IP is not related to gender or marital status, but it is negatively associated with the number of offspring and has a quadratic relationship with age.

For all of the 76 triangle subjects we have one within-task test of IP and three between task tests. We use this data to construct an IP-score for each individual out of 4 (where 0 means behaviour inconsistent with IP in all four tests and 4 means behaviour consistent with IP etc.). A general to specific specification search using an ordered probit model produces the following equation (t-statistics in parentheses) for the index function:

$$Y = 5.46 - 0.15\text{Age} + 0.015\text{Agesq} - 0.19\text{Kids} \quad N = 76$$

$$(3.43) \quad (-1.88) \quad (1.74) \quad (-1.59)$$

Age is the age of the subject, Agesq is the square of their age and Kids is the number of their children. Rejected variables include whether married or not, gender, and years together. Given the sample characteristics, the net effect of age on the IP score is always negative, but there is a minimum at age 47.6 beyond which subjects become more likely to be IP consistent. This result is slightly surprising: it suggests that the youngest people in our sample (who typically also have the fewest children) are the most likely to income pool in individual choice. When we do a separate regression for the 76 wave subjects using data from one within-task IP test and two between task tests, we get very similar results - a negative relationship with the number of children and a minimum point for the negative effects of age at 39 years. Although our sample is relatively small, it seems that some of the factors that might intuitively be associated with a greater degree of income pooling are actually associated with its opposite.²⁰

takes the value 1 if the partner states that the male is looks after the household finances comes close to significance at the 5% 1-tailed level.

²⁰ Interestingly, when we add the partner's IP score as a right-hand side variable it has no additional explanatory power, suggesting that matching is not associative in the propensity to income pool.

5. DISCUSSION

In this paper we have tested experimentally properties of household choice. The results suggest that when couples make choices together, income pooling is a feature of their decisions. When individuals make decisions about their household in the absence of their partners they place approximately equal weight on their own and their partner's payoffs, but their choices are sensitive to some intra-household transfers of income. Nearly always though individuals and couples pay some attention to the correlation of payoffs, so that separability is not a common feature of choices. The choices made by individuals and the choices made by couples are often different, so that unanimity is also rare. Although many choices do respect the Pareto principle, a significant fraction (about 23 per cent of the relevant cases) do not, particularly when one option is clearly riskier than the other.

Taken at face value, the results suggest that no standard economic model of the family has universal support. Two models in particular do not fare well. A simple weighted-average of utility model with no ex post transfers performs particularly badly. Secondly, the strictest version of the unitary model - in which we see a) income pooling and b) all choices, joint and individual, coincide - does not match the data for more than a few individuals.

Recall that market data (e.g. Lundberg et al 1997) has produced conclusions largely hostile to the income pooling assumption. On the other hand our results produce a sizeable amount of behaviour that is compatible with income pooling, even amongst individuals. To reconcile these two facts we note two features of our dataset. First, even when each individual in a household makes decisions consistent with income pooling, those individual decisions are frequently inconsistent with each another, possibly because individual attitudes to risk differ or because individuals have erroneous views of their partner's attitudes to risk. Secondly, many of our households report financial arrangements (e.g. separate budgets) which suggest a significant degree of decentralisation in their decision-making. Such features of households are common. A combination of widespread decentralised decision-making, coupled with imperfect information about partners and differing attitudes to risk could then account for the failure of IP in market data.

References

- Alderman, H., Chiappori, P.A., Haddad, L., Hoddinott, J. and 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.
- Becker, G.S., (1974) A theory of social interactions. *Journal of Political Economy*, 82(6): 1063-1093.
- Bergstrom, T. (1989) A fresh look at the rotten kid theorem and other household mysteries. *Journal of Political Economy*, 97(5): 1138-1159.
- Bone, J., Hey, J. and Suckling, J. (1999) Are Groups More or Less Consistent Than Individuals? *Journal of Risk and Uncertainty*, 8: 63-81.
- Browning, M. and Chiappori, P.A. (1998) Efficient intra-household allocations: a general characterisation and empirical tests. *Econometrica* 66(6): 1241-1278.
- Chen, Z. and Woolley, F. (2001) A Cournot-Nash Model of Family Decision-making. *Economic Journal*, 111: 722-748.
- Corfman, K.P. and Lehmann, D.R. (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, R., Starmer, P., Chris, V. and Sugden, R. (1998) On the validity of the random lottery incentive system. *Experimental Economics*, 1: 115-131.
- Doss, C. (1996) Testing among models of intrahousehold resource allocation. *World Development*, 24(10): 1597-1609.
- Guth, W. and Tietz, R. (1990) Ultimatum Bargaining Behaviour: A Survey and Comparison of Experimental Results. *Journal of Economic Psychology*, 11: 417-449.
- Lundberg, S. and Pollak, R.A. (1993) Separate Spheres Bargaining and the Marriage Market. *Journal of Political Economy*, 101(6): 988-1010.
- Lundberg, S.J., Pollak, R.A. and Wales, T.J. (1997) Do Husbands and Wives Pool Their Resources? Evidence from the United Kingdom Child Benefit. *Journal of Human Resources*, 32(3): 463-480.
- Jones, C.W. (1983) Mobilisation of women's labour for cash crop production: a game theoretic approach. *American Journal of Agricultural Economics*, 65: 1049-2054.
- Manser, M. and M. Brown, (1980), Marriage and Household Decision-Making: A Bargaining Analysis', *International Economic Review.*, 21 (1), 31-44..
- Pahl, J. (1990) Household spending, personal spending and the control of money in marriage. *Sociology*, 24(1): 119-138.
- Peters, H.E., Ünür, A.S., Clark, J. and Schulze, W.D. (forthcoming) Free-Riding and the Provision of Public Goods in the Family: An Experimental Test of the Rotten Kid Theorem. *International Economic Review*.
- Phipps, S.A. and Burton, P.S. (1998) What's Mine is yours? The influence of male and female

- incomes on patterns of household expenditure. *Economica*, 65: 599-613.
- Rothschild, M. and Stiglitz, J.E. (1970) Increasing risk I: A definition. *Journal of Economic Theory*, 2: 225-243.
- Strauss, J., Mbabu, G. and Beegle, K. (2000) Intra-household allocations: a review of theories and empirical evidence. *Journal of African Economics*, 83-143.
- Treas, J. (1993) Money in the Bank: Transaction Costs and the Economic Organization of Marriage. *American Sociological Review*, 58: 723-734.
- Udry, C. (1996) Gender, Agricultural Productivity and the Theory of the Household. *Journal of Political Economy*: 104/5.
- Udry, C. and Duflo, E. (2001) Intrahousehold Resource Allocation in Côte d'Ivoire: Social Norms, Separate Accounts and Consumption Choices. Yale University *mimeo*, July.
- Warr, P.G. (1983) The private provision of a public good is independent of the distribution of income. *Economic Letters*, 13: 207-11.
- Woolley, F. (2000) Control over Money in Marriage. *Carleton Economic Papers*, 00-07.

Appendix A: The tasks

		Lottery 1				Lottery 2			
		Triangle		Wave		Triangle		Wave	
		£20	£40	£20	£40	£20	£40	£20	£40
Group 1	Set 1: T1, W1, J1	1-100	-	-	-	21-70	71-100	-	-
	Set 2: T2, W2, J2	21-100	-	1-20	-	21-100	-	71-100	-
	W8	-	-	1-100	-	-	-	21-70	71-100
Group 2	Set 3: T3, W3, J3	1-100	-	-	-	-	-	-	41-100
	W9, J8	1-50	-	51-100	-	41-100	-	41-100	-
	T8	1-100	-	-	-	-	41-100	-	-
Group 3	Set 4: T4, W4, J4	31-100	-	1-100	-	21-70	71-100	-	21-100
	T9, J9	1-100	-	31-100	-	-	21-100	21-70	71-100
Group 4	Set 5: T5, W5, J5	-	-	51-100	-	-	-	-	71-100
	T10	51-100	-	-	-	-	71-100	-	-
Group 5	Set 6: T6, W6, J6	1-100	-	1-100	-	1-100	-	21-70	71-100
	W10	1-100	-	1-100	-	21-70	71-100	1-100	-
Group 6	Set 7: T7, W7, J7	-	-	1-100	-	-	1-70	-	-
Group 7	T11, J10	51-100	-	1-50	-	1-100	-	-	71-100
	W11	31-100	-	1-30	-	-	71-100	1-100	-
Other	T12	51-100	-	31-100	-	1-100	-	-	71-100

W12	31- 100	-	51- 100	-	-	71- 100	1-100	-
T13, W13, J11	1-50	-	51- 100	-	1-50	51- 100	51- 100	-
J12	71- 100	-	1-70	-	1-50	-	1-70	-
W14	1-70	-	21-70	71- 100	71- 100	-	1-70	-
W15	1-100	-	1-100	-	1-100	-	-	41- 100
J13	-	71- 100	1-70	-	1-60	61- 100	-	-
J14	-	-	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.

Appendix B:

Subject pair__Partner_Δ_

Section 1

In this section you will face a series of twelve different choices involving options like example 1 shown below.

Each choice will involve two options, labelled A and B.

For each question your task is to choose the option you would prefer to have.

At the end of the experiment, one question number from 1-36 will be drawn at random.

If one of these questions is selected at the end of the experiment you will play out the option chosen by you.

You will be asked to draw a number from 1-100 from a bag containing 100 discs.

You will receive your prize corresponding to that number and your partner will receive their prize corresponding to the same number.

Your choices are confidential - we will not reveal them to your partner.

The options you will face are similar to the one shown in Example 1:

Example 1.

For numbers:	1-20	21-70	71-100
You receive	£0	£0	£40
Your partner receives	£0	£20	£40

The row beginning 'For numbers' has three ranges of numbers, corresponding to the numbers in the bag. Below each range is a prize for you and then a prize for your partner.

For instance, under the range 21-70 you can see a prize of £0 for you and a prize of £20 for your partner. These are the prizes you would get if you had this option and you pulled a number between 21 and 70 from out of the bag.

(please turn over)

Example 2 shows choices as you will see them. In each case you must choose ONE option by ticking the appropriate box.

Example 2:

Option A		
For numbers:	1-60	61-100
You receive	£20	£20
Your partner receives	£0	£20

Option B		
For numbers:	1-20	21-70
You receive	£0	£40
Your partner receives	£0	£20

I choose (*tick one*): Option A ☐

Option B ☐

For instance, suppose you chose Option B by ticking the relevant box, then if the number 25 was pulled from the bag at the end of the experiment, you would receive nothing and your partner would get £20.

Check your understanding.

Suppose that in the example question you had chosen Option B and that this question was picked at the end of the experiment.

What would you receive if the number 80 was picked from the bag? (*tick one*) 0 ☐ £20 ☐ £40 ☐

What would your partner receive if the number 20 was picked from the bag? (*tick one*) £0 ☐ £20 ☐ £40 ☐

To sum up, for each question your task is to choose the option you prefer.